

The Trademark Reporter®



The Law Journal of the International Trademark Association

ARTIFICIAL INTELLIGENCE ISSUE

AI and the “Death of Trademark”

Michael Grynberg

Trademarks in an Algorithmic World

Christine Haight Farley

Trademark Confusion Revealed: An Empirical Analysis

Daryl Lim

Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models

Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E. Ho

Trademark Search, Artificial Intelligence, and the Role of the Private Sector

Sonia K. Katyal and Aniket Kesari

Generating Rembrandt: Artificial Intelligence, Copyright, and Accountability in the 3A Era—the Human-Like Authors Are Already Here—a New Model

Shlomit Yanisky-Ravid

Copyright Thickness, Thinness, and a *Mannion* Test for Images Produced by Generative Artificial Intelligence Applications

Molly Torsen Stech

INTERNATIONAL TRADEMARK ASSOCIATION

675 Third Avenue, New York, NY 10017-5704

Telephone: +1 (212) 642-1700

email: wknnox@inta.org

OFFICERS OF THE ASSOCIATION

DANA NORTHCOTT	<i>President</i>
ELISABETH BRADLEY	<i>President-Elect</i>
DEBORAH HAMPTON	<i>Vice President</i>
MARLOU VAN DE BRAAK	<i>Vice President</i>
SERGIO BARRAGAN	<i>Treasurer</i>
DEBRA HUGHES	<i>Secretary</i>
GUSTAVO GIAY	<i>Counsel</i>
ETIENNE SANZ DE ACEDO	<i>Chief Executive Officer</i>

The Trademark Reporter Committee

EDITORIAL BOARD

EDITOR-IN-CHIEF, CHAIR

PAMELA CHESTEK

STAFF EDITOR-IN-CHIEF

WILLARD KNOX

VICE CHAIR

ANDREW J. GRAY

Senior Editors

KAREN L. ELBURG
LESLEY GROSSBERG

VERENA VON BOMHARD
DÉSIRÉE FIELDS

BARRY HORWITZ

Director of Legal Resources

LIZ HANELLIN

Senior Staff Editor

BEVERLY HARRIS

Senior Staff Editor

ELIZABETH VENTURO

Editor-in-Chief, Practice Guides and Legal Resources

ROSEMARY DESTEFANO

Composer

BARBARA MOTTER

Editors

SAFIR ANAND
JUSTIN ANDERSON
RICARDO ENRIQUE ANTEQUERA
DOUG BANIA
DANIEL R. BERESKIN
IAN BETTS
KATHLEEN E. BRENNAN
LANNING G. BRYER
ANDREA CAPPAL
EDWARD CHATTERTON
VIVIAN CHENG
STEPHANIE CHILD
THEODORE H. DAVIS JR.
NOAH DRAKE
HENRY DU
SAMUEL EICHNER
CLAIRE ELLER
J. KEVIN FEE
DYAN FINGUERRA-DUCHARME
MATHILDE P. FLORENSON
LAURA GAO
ALEX GARENS
NAOMI JANE GRAY
TAYLAR GREEN

ULRIKE GRUEBLER
JORDI GÜELL
ÖZGÜL GÜNGÖR
WINSLOW HALL
CONNOR HANSEN
ASHNEET HANSPAL
JOHN HEINBOCKEL
CHRISTOPHER HOOLEHAN
BRUCE ISAACSON
MANAVI JAIN
PETER KAROL
ELISABETH KASZNAR FEKETE
MOHSIN KIDWAI
STEVEN LAURIDSEN
SCOTT LEBSON
BARBARA ANGELA LEITAO-JOHNSON
REBECCA S. LESSARD
NEIN-CHIN LIAO
ILARIA MAGGIONI
SHABNAM MALEK
J. THOMAS MCCARTHY
RANDY MICHELS
ARANTXA MONROY

ERIC R. MORAN
ABHILASHA NIROOLA
CHRISTOPHER PARENT
NEAL ROBERT PLATT
NATHAN RANNS
PARTEL-ERIK ROUK
ANDREA RUSH
LEIGHA R. SANTORO
TOM DANIEL SCOURFIELD
GABRIEL SILVEIRA SACRAMENTO
MARIA SINATRA
MEGAN SNEED
BRIAN SOWERS
CARL J. SPAGNUOLO
KELU SULLIVAN
DONALD THOMPSON
JODI-ANN TILLMAN
JULIE TORTORA
SANAYA FAROKH VACHHA
MARTIN VIEFHUES
JOHN L. WELCH
BRYAN WHEELOCK
JASON W. WHITNEY
LAILA SENGUPTA WOLFGRAM

Advisory Board

MILES J. ALEXANDER
WILLIAM M. BORCHARD
LANNING G. BRYER
JESSICA ELLIOTT CARDON
SANDRA EDELMAN
ANTHONY L. FLETCHER

ROBERT M. KUNSTADT
THEODORE C. MAX
KATHLEEN MCCARTHY
GLENN MITCHELL
JONATHAN MOSKIN
JOHN B. PEGRAM

ROBERT L. RASKOPF
PASQUALE A. RAZZANO
PIER LUIGI RONCAGLIA
HOWARD SHIRE
JERRE B. SWANN, SR.
RAFFI V. ZEROUNIAN

The views expressed in *The Trademark Reporter* (TMR) do not necessarily reflect those of the International Trademark Association (INTA). To fulfill its mission of delivering cutting-edge scholarship on trademarks, brands, and related intellectual property to its readers, the TMR sources content reflecting a diversity of viewpoints; the views expressed in any given article, commentary, or book review are those of the individual authors.

The TMR (ISSN 0041-056X) is published electronically six times a year by INTA, 675 Third Avenue, New York, NY 10017-5704 USA. INTA, the INTA logo, INTERNATIONAL TRADEMARK ASSOCIATION, THE TRADEMARK REPORTER, and inta.org are trademarks, service marks, and/or registered trademarks of the International Trademark Association in the United States and certain other jurisdictions.

The Trademark Reporter®

(USPS 636-080)

Vol. 114

November–December, 2024

No. 6

TABLE OF CONTENTS

Editor’s Note: Theme Issue—Artificial Intelligence..... 691

ARTICLES

AI and the “Death of Trademark” Michael Grynberg.....	695
Trademarks in an Algorithmic World Christine Haight Farley	735
Trademark Confusion Revealed: An Empirical Analysis Daryl Lim	799
Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E. Ho	880
Trademark Search, Artificial Intelligence, and the Role of the Private Sector Sonia K. Katyal and Aniket Kesari.....	910
Generating Rembrandt: Artificial Intelligence, Copyright, and Accountability in the 3A Era—the Human-Like Authors Are Already Here—a New Model Shlomit Yanisky-Ravid.....	998
Copyright Thickness, Thinness, and a <i>Mannion</i> Test for Images Produced by Generative Artificial Intelligence Applications Molly Torsen Stech.....	1066

The Trademark Reporter®

EDITOR'S NOTE:

THEME ISSUE—ARTIFICIAL INTELLIGENCE

We are pleased to devote this end-of-year issue of *The Trademark Reporter* (TMR) to the realm of artificial intelligence (“AI”). Whether you believe that AI is mostly hype or as seismic a change as the smartphone, there is no doubt that it will evolve and improve, fundamentally changing both our professional and private lives in ways we can only begin to imagine.

In this issue we are looking, narrowly and broadly, at what impact AI may have on trademarks—the choice of trademarks, the effect of AI on consumer choice, and its impact on jurisprudence. We have collected articles from preeminent authors that describe both where we are now and what we might expect the future to bring.

FOUNDATIONAL TRADEMARK THEORY IN AN AI WORLD

We start with Michael Grynberg’s “AI and the ‘Death of Trademark.’”¹ Published in 2020, before the introduction of ChatGPT by its creator, Open AI, in November 2022,² Grynberg’s article imagines an AI assistant—with abilities that, at the time, Grynberg considered hypothetical but that are now fully realized in today’s AI models—that will be able to make choices about products based on your preferences and sort through the overwhelming amount of data about them to find the best choice for you. Grynberg posits that trademarks will take a lesser role in product selection, even perhaps becoming superfluous. The author points out that we have already entered this world through Amazon.com, where consumers rely on the search and suggestion algorithms and product ratings rather than nonsensical, unpronounceable trademarks when making product choices.

Christine Haight Farley takes the next step, pointing out in her 2023 article “Trademarks in an Algorithmic World”³ that the very foundation of modern trademark law, the “search costs theory,” may

¹ 108 Ky. L.J. 199 (2020).

² See <https://openai.com/index/chatgpt/>.

³ 98 Wash. L. Rev. 1123 (2023).

no longer apply when we have AI agents making product choices. Farley points out that trademark doctrine relies almost entirely on the Chicago School of Economics theory that trademarks create an efficient market by reducing search costs. But this justification for protecting trademarks may no longer be valid in a world where AI is far superior to a trademark at identifying products with the desired qualities. However, trademarks have value other than aiding in search, namely, the brand as product per se, or, as Farley puts it, the “manufactured desire” for a product. Farley argues that, if trademarks are still to be protected, the jurisprudence must find different policy justifications for how and why the law will protect the brand producer when search efficiency, as incomplete a theory as it has been for some time, may be entirely irrelevant.

THE USE OF AI IN LEGAL DECISION-MAKING

We next move from trademark theory to the use of AI in decision-making. In “Trademark Confusion Revealed: An Empirical Analysis,”⁴ Daryl Lim analyzes courts of appeals’ decisions on the likelihood of trademark confusion and how the various factors are applied by the courts. Lim considers the possibility of using empirical analysis to train AI models that would allow courts to reach more consistent and accurate results and that would, in turn, allow trademark counselors to better predict likely outcomes before ever reaching the courts.

But if one is inclined to turn our legal decision-making over to our AI overlords, authors Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E. Ho educate us on the inherent flaws in large language models. In “Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models,”⁵ the authors analyze the extent of legal hallucinations by different Large Language Models, with the degree of hallucinations when evaluating district court cases as high as 58% (in the ChatGPT 4 model⁶) to 88% (in the Llama 2 model created by Meta⁷). The authors point out that hallucinations are inevitable as the AI model tries to maintain fidelity to the training corpus, the user’s prompt (which may assume incorrect information), and the law itself, three areas that may be in conflict. And hallucinations are not always bad: for example, a hallucination may help the litigator develop new theories—thus, useful to the litigator, but perhaps something we do not want a court to readily adopt. Incorporating AI into legal decision-making will be challenging to get right.

⁴ 71 Am. U. L. Rev. 1285, 1289 (2022).

⁵ 16 J. Legal Analysis 64 (2024).

⁶ See <https://openai.com/index/gpt-4/>.

⁷ See <https://www.llama.com/llama2/>.

CREATION OF RIGHTS

Finally, we look at the role of AI in the adoption of trademarks. Sonia K. Katyal and Aniket Kesari, the authors of “Trademark Search, Artificial Intelligence, and the Role of the Private Sector,”⁸ explore the realm of private trademark search providers, already using AI, and their effect on trademark selection. In 2019 and 2020, the authors conducted experiments comparing the performance of different AI-powered trademark search engines at identifying potential likelihood of confusion refusals. They propose that not just consumers but also trademark owners should be considered relevant economic actors in the trademark ecosystem, with a significant interest in selecting trademarks that will not be rejected by trademark offices.

We then begin to think more broadly about the protection of trademarks based on their parallel existence as creative works. Copyright is often used as a secondary means of protection of trademark rights when a logo is sufficiently original and creative to merit copyright protection. However, currently the U.S. Copyright Office is parsimonious in its willingness to register copyrights in logos, and now is also reserved in its grant of copyright protection for AI-generated works. This confluence of Copyright Office doctrine does not bode well for trademarks, even where the use of AI may have been modest. In 2017, Shlomit Yanisky-Ravid authored “Generating Rembrandt: Artificial Intelligence, Copyright, and Accountability in the 3A Era—The Human-like Authors are Already Here—A New Model,”⁹ in which the author suggests that we adopt a new ownership model for AI-generated works based on the existing work-made-for-hire laws.

And we end the issue with an article by Molly Torsen Stech, “Copyright Thickness, Thinness, and a *Mannion* Test for Images Produced by Generative Artificial Intelligence Applications.”¹⁰ Stech also searches for an answer to the conundrum of copyright ownership of works generated through the use of, or assistance from, AI. Stech suggests taking a cue from the jurisprudence of the copyrightability of photographic works as considered in *Mannion v. Coors Brewing Co.*¹¹ Photography is a medium commonly compared to AI because of the speed and rapidity with which works can be created, and the role of happenstance (or luck) in a work’s creation. The court in *Mannion* opined that copyrightability of photographic works can be determined by considering the work’s rendition, its timing, and the photographer’s role in creation of the subject. Stech

⁸ 35 Berkeley Tech. L.J. 501 (2020).

⁹ 2017 Mich. St. L. Rev. 659 (2017).

¹⁰ 2024 B.C. Intell. Prop. & Tech. F. 1 (2024).

¹¹ 377 F. Supp. 2d 444 (S.D.N.Y. 2005).

suggests applying a modified *Mannion* test to AI-created works to ascertain an appropriate dividing line between uncopyrightable, machine-generated works and ones where the human intervention is sufficiently high to reward the creator with copyright protection.

We hope you find this issue of the TMR interesting and illuminating, as we all begin to explore this new technology and begin to learn and understand how it will shape, in small ways or big ways, our profession and our practice.

Pamela Chestek
Editor-in-Chief

Kentucky Law Journal

VOLUME 108

2019–2020

NUMBER 2

ARTICLES

AI AND THE “DEATH OF TRADEMARK”

*Michael Grynberg*¹

TABLE OF CONTENTS

TABLE OF CONTENTS	199
INTRODUCTION.....	200
I. AI SHOPPERS.....	201
A. <i>Machine learning</i>	201
B. <i>The digital shopper</i>	202
C. <i>Other than that, Mrs. Lincoln, how was the play?</i>	203
D. <i>AIs and trademarks</i>	204
II. THE “END” OF TRADEMARK LAW?	205
A. <i>The role of trademarks</i>	205
B. <i>What’s left?</i>	210
i. Maintaining distributive fairness	210
ii. Preserving a residual core	211
iii. Everything is false advertising	211
iv. Everything is unfair competition and the problem of prestige goods.....	212
C. <i>What’s gained?</i>	213
i. Legalizing competition	213
ii. Opening communications channels	214
iii. Protecting free expression	215
D. The “death of trademark” in a bespoke world.....	215
III. KNOWLEDGE ONLINE AND TRADEMARKS.....	216
A. <i>Knowledge Online and Trademarks</i>	216
B. <i>Trademarks and the stability of meaning</i>	219
C. <i>Stability of meaning online: Two views</i>	221
i. The internet and the collapse of market barriers	221
1. <i>The story</i>	221
2. <i>The history</i>	222
ii. The internet and the retreat of trademark authority	223
1. <i>The story</i>	223
2. <i>The history</i>	224
IV. NASCENT AIs.....	225
A. <i>Amazon and trademark’s online retreat</i>	226
B. <i>One brand to rule them all?</i>	229
C. <i>Controlling avatars</i>	232
V. THE LARGER LIMITS OF TRADEMARK	234
CONCLUSION.....	238

¹ Professor of Law, DePaul University College of Law. Comments welcome at mgrynber@depaul.edu. My thanks to the participants at the 8th Annual Trademark Works in Progress Colloquium at the American University Washington College of Law, Robert Brauneis, Christine Haight Farley, Yvette Joy Liebesman, Dustin Marlan, Alexandra Roberts, Zvi Rosen, Matthew Sipe, and Julie Cromer Young, for their comments and feedback. I’m also grateful for comments received at the Works in Progress Intellectual Property (WIPIP) Colloquium at the University of Houston Law Center and from my colleagues at DePaul University College of Law.

INTRODUCTION

How might improvements in artificial intelligence (AI) technology affect trademark law? This Article approaches the question by imagining trademark law in a world in which we can fully outsource our consuming decisions to AIs that know our preferences better than we do. Leaving aside whether the technology is possible, this thought experiment tells us something about today's trademark doctrine and its response to changes in online technology and culture.

Part I imagines a hypothetical AI capable of assuming responsibility for our lives as consumers. Part II argues that a sufficiently sophisticated technology would render trademarks superfluous in many cases. Trademarks function by simplifying information. We use them to stand in for a broad range of (sometimes contradictory) data assembled from a variety of sources. Because human cognition is limited, the ability of trademarks to serve as a shortcut is valuable, but it is a second-best solution. With unlimited time or enhanced capacities, we would be better able to find optimal products without relying on the simple information signals offered by trademarks. As it is, sifting through all the available data is not a wise use of our limited attention.

But the hypothetical AI is not similarly limited, and a sufficiently advanced AI "shopper" would exist to analyze the context that trademarks allow us to ignore. The role of trademarks in such a world is more limited—and consequently requires less protection—than what we see today.

Part III explains that the hypothetical AI also illuminates a tension between trademark law and the consumption of knowledge online. Before the internet, the relative scarcity of "space" for information—be it on library shelves, newspaper pages, or television channels—conferred authority on those—be they librarians, editors, or programmers—able to curate it. Not so online. Comparatively speaking, there is room enough for practically anything. We therefore rely on filtration, rather than curation, to find information; our filters leave the rejected data available for others to find and use as they see fit.² This facilitates the formation of communities that have mutually irreconcilable conceptions of truth. These disagreements could extend to trademark meaning, but the current model of trademark information reflects, to a large extent, the scarcity model that has proven ill-adapted to life online.

Of course, the hypothetical AI does not, and may never, exist. As Part IV explains, however, we can see its forerunners in web platforms like Amazon and Facebook. They are already changing trademark doctrine, and they illustrate why trademarks may be less important in the future. And of course, these technologies raise any number of troubling questions, but they are not the sort that trademark law is designed to address.

² DAVID WEINBERGER, *TOO BIG TO KNOW* 11 (2011).

I. AI SHOPPERS

A. Machine learning

Computers now surpass humans at any number of skills associated with cognition. Chess and checkers fell a long time ago,³ and computers now reign supreme at Go, a game once thought to be too complex for AI.⁴ Even poker, with the human element of bluff seemingly built into the rules of the game, is not immune.⁵ And a variety of tasks once thought to be the exclusive domain of trained professionals are ever increasingly open to automation.⁶

Many recent advances are popularly associated with machine learning. Rather than program a computer with predetermined algorithms that channel the machine’s processing power, this approach leaves it to computers to sort out methods for themselves.⁷ Computers can find patterns in the data they receive (or generate) and then apply derived rules to the assigned task, refining them iteratively.⁸ As a result, a computer may teach itself to, say, defeat a human at the game of Go, but its internal rules for selecting a good move may diverge from those used by human professionals.⁹ The results work, but the steps followed by the machine are opaque

³ Alexis C. Madrigal, *How Checkers Was Solved*, ATLANTIC (July 19, 2017), <https://www.theatlantic.com/technology/archive/2017/07/marion-tinsley-checkers/534111/> [<https://perma.cc/4FMA-RU27>]; see Olivia Solon, *Oh the Humanity! Poker Computer Trounces Humans in Big Step for AI*, GUARDIAN (Jan. 30, 2017, 17:00 EST), <https://www.theguardian.com/technology/2017/jan/30/libratus-poker-artificial-intelligence-professional-human-players-competition> [<https://perma.cc/ADD7-PHAV>].

⁴ See Christopher Moyer, *How Google’s AlphaGo Beat a Go World Champion*, ATLANTIC (Mar. 28, 2016), <https://www.theatlantic.com/technology/archive/2016/03/the-invisible-opponent/475611/> [<https://perma.cc/7XHR-KQKT>].

⁵ Solon, *supra* note 3. Similarly, an AI system prevailed against top players in the war strategy game StarCraft. Kelsey Piper, *StarCraft is a Deep, Complicated War Strategy Game. Google’s AlphaStar AI Crushed It.*, VOX (Jan. 24, 2019, 7:04 PM), <https://www.vox.com/future-perfect/2019/1/24/18196177/ai-artificial-intelligence-google-deepmind-starcraft-game> [<https://perma.cc/W4US-SA82>].

⁶ See, e.g., Steve Lohr, *A.I. is Doing Legal Work. But It Won’t Replace Lawyers, Yet.*, N.Y. TIMES (Mar. 19, 2017), <https://www.nytimes.com/2017/03/19/technology/lawyers-artificial-intelligence.html> [<https://perma.cc/C34X-ANSQ>]; Jessica Stillman, *An A.I. Just Outperformed 20 Top Lawyers (and the Lawyers Were Happy)*, INC. (Nov. 9, 2018), <https://www.inc.com/jessica-stillman/an-ai-just-outperformed-20-top-lawyers-and-lawyers-were-happy.html> [<https://perma.cc/GR9H-MJAC>].

⁷ See David Silver et al., *A General Reinforcement Learning Algorithm that Masters Chess, Shogi, and Go through Self-play*, SCI., Dec. 7, 2018, at 1140, 1140, <https://science.sciencemag.org/content/sci/362/6419/1140.full.pdf> [<https://perma.cc/M6NS-EU9A>] (“Our results demonstrate that a general-purpose reinforcement learning algorithm can learn, tabula rasa—without domain-specific human knowledge or data, as evidenced by the same algorithm succeeding in multiple domains—superhuman performance across multiple challenging games.”).

⁸ See *id.* (“A long-standing ambition of artificial intelligence has been to create programs that can instead learn for themselves from first principles. Recently, the AlphaGo Zero algorithm achieved superhuman performance in the game of Go, by representing Go knowledge using deep convolutional neural networks, trained solely by reinforcement learning from games of self-play.” (footnotes omitted)).

⁹ As when AlphaGo defeated the human champion Lee Sedol in the game of Go:

With the 37th move in the match’s second game, AlphaGo landed a surprise on the right-hand side of the 19-by-19 board that flummoxed even the world’s best Go players, including Lee Sedol. “That’s a very strange move,” said one commentator, himself a nine dan

to humans.¹⁰ Indeed, in some cases computers need only the rules of the game, and they can take it from there.¹¹

Though the precise line between hype and actual potential is unclear, machine learning technology can be put to any number of uses. Google famously used it to change its translation software, discarding years of effort that focused on algorithmic translation rules and dictionary databases.¹² The company found better results when it fed a program an extensive library of works and translations, letting the program discern translation rules of its own.¹³

It seems reasonable to assume that today's high-end AI technology will continue to develop and eventually find its way into consumer products, deployable to a range of ends. Suppose we use it to manage our purchasing decisions?

B. *The digital shopper*

Various forms of AI already mediate our lives as consumers,¹⁴ but thinking about potential future developments may provide interesting insights into trademark law. Imagine—no claims are made that this thought experiment is necessarily plausible in all its details—an AI far surpassing today's rudimentary digital assistance tools. This hypothetical digital shopper could fully manage your purchasing choices, as the AI can be delegated the task of researching, evaluating, and purchasing goods and services on your behalf. Your digital personal shopper will be trained in your preferences and possessions, supplemented as necessary by further input, and it will then extrapolate the nature of goods and services that are likely to please. It may err from time to time—just as you do—but soon enough it will know you well enough to outperform you in predicting what you'll like. The advantage will not simply be one of time saved—in which the AI delivers a second-best choice that satisfies because you could skip the effort of shopping—the machine will be better at figuring

Go player, the highest rank there is. "I thought it was a mistake," said the other. Lee Sedol, after leaving the match room, took nearly fifteen minutes to formulate a response. . . .

Indeed, the move turned the course of the game. AlphaGo went on to win Game Two, and at the post-game press conference, Lee Sedol was in shock. "Yesterday, I was surprised," he said through an interpreter, referring to his loss in Game One. "But today I am speechless. If you look at the way the game was played, I admit, it was a very clear loss on my part. From the very beginning of the game, there was not a moment in time when I felt that I was leading."

It was a heartbreaking moment. But at the same time, those of us who watched the match inside Seoul's Four Seasons hotel could feel the beauty of that one move. . . .

Cade Metz, *In Two Moves, AlphaGo and Lee Sedol Redefined the Future*, WIRED (Mar. 16, 2016, 7:00 AM), <https://www.wired.com/2016/03/two-moves-alpha-go-lee-sedol-redefined-future/> [<https://perma.cc/2QE9-K7J8>].

¹⁰ See Will Knight, *The Dark Secret at the Heart of AI*, MIT TECH. REV. (Apr. 11, 2017), <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/> [<https://perma.cc/X8NH-6NGQ>].

¹¹ See, e.g., Silver et al., *supra* note 7, at 1140.

¹² See Gideon Lewis-Kraus, *The Great A.I. Awakening*, N.Y. TIMES (Dec. 14, 2016), <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html> [<https://perma.cc/79R8-VVHP>].

¹³ See *id.*

¹⁴ See *infra* Part IV.

out what you’ll like. The AI’s ability to sift through the flood of available data in tireless search of optimal results will further enhance its advantage.

Even though the AI knows your desires better than you do, you will not understand why. Maybe the suggestion to eat at the new creperie owes its origin to a political donation, a song on your playlist, your hometown, or some combination of these or other details. Who knows why? The AI sees a pattern, and it works.¹⁵

Imagine further that this experience is so typical as to be banal. Everyone takes for granted that their consumer avatars are as much a part of the fabric of life as smartphones are today. They take care of business and are given thought only on the rare occasion that something goes wrong, and the AI needs the gentle guidance of a thumbs down button (or its future equivalent).

C. Other than that, Mrs. Lincoln, how was the play?

There are any number of possible objections to the framing of the hypothetical. Some concern the various reasons that it may be unrealistic. For example, the hypothesized AI relies on an ability to understand language that contemporary AI lack and may never achieve. The response is simply that this is a thought experiment, not a prediction. As we say to our students, don’t fight the hypothetical.¹⁶

More importantly, the above account of a personal assistant AI will likely strike many as dystopic, perhaps because of the privacy implications for the user or for society at large—insofar as the technology might require considerable personal and third-party data to make effective predictions on behalf of its user.¹⁷ That data must be collected and assembled, creating a privacy security risk. Worse, getting individuals to generate and reveal the data necessary to feed the machine may require an unacceptable level of individual manipulation. That is, devices may be engineered to prod people into providing the data that may then be the basis of AI learning on a scale beyond what already generates alarm today.¹⁸ This is independent of other

¹⁵ As when Netflix recommends things I have already seen because I recently watched something else, but with greater accuracy and utility.

¹⁶ Though to this particular objection, perhaps full linguistic comprehension is unnecessary. If IBM Watson can understand Jeopardy questions well enough to answer them, it is not clear that non-comprehending, but advanced, AI would be incapable of shopping for goods.

¹⁷ Even then, that data may be too retrospective to provide useful predictions for novel situations. See CATHY O’NEIL, *WEAPONS OF MATH DESTRUCTION: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY* 204 (2016) (“Big Data processes codify the past. They do not invent the future.”).

¹⁸ See SHOSHANA ZUBOFF, *THE AGE OF SURVEILLANCE CAPITALISM: THE FIGHT FOR A HUMAN FUTURE AT THE NEW FRONTIER OF POWER* 284–85 (2019) (“The new toolmakers do not intend to rob you of your inner life, only to surveil and exploit it. All they ask is to know more about you than you know about yourself.”); *id.* at 241 (“All that is moist and alive must hand over its facts. There can be no shadow, no darkness. The unknown is intolerable.”). Zuboff sees the development of digital assistants as part of this process. See *id.* at 255–60. The dynamic she describes applies, however, to a wide range of technology. In her account, the drive is to render an ever-greater amount of data as fodder for the prediction markets that depend on it. See, e.g., *id.* at 236–38 (describing data collected by Nest thermostat and mattress companies and the practical difficulty of preventing its collection). Furthermore, using devices engineered to encourage use creates potential psychological harm independent of the underlying goal. See generally ADAM ALTER, *IRRESISTIBLE: THE RISE OF ADDICTIVE TECHNOLOGY AND THE*

manipulation possibilities, such as the prospect that the AI might nudge consumers into particular purchases or other decisions.¹⁹

As important as these considerations are, they are largely outside the scope of this Article,²⁰ which is interested in the hypothetical AI as a tool for examining trademark law. It does not advocate for its development, nor does it take a position on how or whether consumer AI technology should develop more generally. My bracketing off certain policy questions is not to deny the prospect that the costs of an AI shopper might outweigh the benefits.

D. AIs and trademarks

Thinking about the implications for trademark law requires that we refine the hypothetical further. First, assume that an AI will encounter trademarks in ways analogous to life online today. Today, one encounters marks in content, including advertising, seller web pages, product reviews by professionals (e.g., a *New York Times* restaurant review) and fellow consumers (e.g., Yelp or Amazon.com reviews), blog and social media references, appearances in videos, etc. Trademarks also appear in information locators (e.g., URLs) and metadata. In the future, trademarks presumably will continue to be used for a variety of purposes, and some of these uses may be deceptive or confusing perhaps by intention, perhaps not. And as today, the trademarks will be surrounded by context that may limit or exacerbate the potential for confusion.

But AI shoppers will have the capability and patience to gather and process far more distinguishing context than humans. A trademark may just be one input among many considered by the AI, and these additional inputs will enable the AI to understand more easily than a human the meaning of a mark in a given context and its relevance, if any, to the AI's goal. This purpose is *not* to distinguish trademarks, but rather to satisfy the preferences of the human the AI serves (unless one of these preferences *is* for authentic trademarks).²¹

The next question then concerns the AI's capabilities. Can it be confused or gamed? We can imagine a range of possibilities here including, at the far end, the

BUSINESS OF KEEPING US HOOKED (2017) (exploring the rise of behavioral addiction by users of digital technology); *id.* at 93–233 (describing “ingredients” of behavioral addiction and how they appear in digital technologies).

¹⁹ See, e.g., Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995, 999 (2014) (“[T]he digitization of commerce dramatically alters the capacity of firms to influence consumers at a personal level.”). This is tied to the question of whom the machine serves, for the right kind of assistant might help us resist external manipulation. See *infra* note 177–178 and accompanying text.

²⁰ A growing body of scholarship addresses these and other implications of AI involvement in consumer decisions. See, e.g., Rory Van Loo, *Digital Market Perfection*, 117 MICH. L. REV. 815, 815 (2019) (“Advancing consumer welfare in the automated era requires not just consumer protection, but digital intermediary protection”); *id.* at 817–18 (collecting examples of scholarship on digital intermediaries); Calo, *supra* note 19. I return to some of these issues below, but with a focus on their interplay with trademark law. Though I do return to them below. See *infra* Parts IV–V.

²¹ As noted above, for now, I am avoiding the question of what the “true” preferences—if such exist—are of any given consumer. See *supra* note 19. I take consumer preferences to be constructed out of a mix of endogenous and exogenous inputs, and I assume they will continue to be so in the future. I am not speculating how a world of ubiquitous, advanced AI might change the mix. I return to the issue below.

prospect of an “omniscient” AI that cannot be misdirected by a false use of a trademark. But even short of that, we might picture lesser AIs with superhuman resistance to deception. An AI that outperforms humans generally may still deliver the occasional “wrong” result due to external manipulation, but its capacity to learn should make these errors unlikely to recur.²²

As this is a thought experiment, we can imagine any level of AI proficiency. We begin at one end of the spectrum. What are the implications of AIs so powerful that confusion as currently contemplated by trademark law becomes, for all practical purposes, impossible? And how does this speculated endpoint of technological development relate to issues in contemporary trademark law? Parts II and III suggest ways that the hypothetical AI may challenge fundamental premises of trademark law.²³ As it is, current technology reflects these challenges in an early form. In other words, as explored in Part IV, many underlying assumptions of trademark law are already undermined by improving digital technology.²⁴

II. THE “END” OF TRADEMARK LAW?

A sufficiently powerful AI could upend trademark law by radically reducing the relevance of trademarks. Today, trademarks simplify information by removing context. But the hypothesized AI’s advantage lies in its ability to sort through context. This capability reduces the importance of maintaining stable trademark meanings, which exist in large part to allow cognitively limited humans to simplify decisions by *ignoring* the context in which trademarks appear. But if our tools were able to put that context to use, trademarks would be left with less to do.

A. *The role of trademarks*

To see why, consider how trademarks function. They are defined by their ability to help consumers identify and distinguish goods and services in the marketplace.²⁵ Once in place, a valid mark conveys information via simplification, letting consumers make assumptions without asking deeper questions about marketplace context. So a buyer may ask for a COKE without wondering how the particular seller defines the term. COKE—vagaries of corporate ownership and licensing aside—designates a “single” source. Likewise, consumers may assume the relevance of their past experience with a mark.²⁶ For instance, eating at a MCDONALD’s in Connecticut provides relevant data about one in Oregon.

²² To be sure, consumer-serving AIs may end up in an arms race with parallel technologies designed to deceive, resulting in an equilibrium that still allows for deception in the marketplace.

²³ See *infra* Parts II–III.

²⁴ See *infra* Part IV.

²⁵ 15 U.S.C. § 1127 (2018).

²⁶ See *Qualitex Co. v. Jacobson Prods. Co.*, 514 U.S. 159, 163–64 (1995) (“In principle, trademark law, by preventing others from copying a source-identifying mark, ‘reduce[s] the customer’s costs of shopping and making purchasing decisions,’ for it quickly and easily assures a potential customer that *this* item—the item with this mark—is made by the same producer as other similarly marked items that he or

Trademarks also allow sellers to assume the expense of assembling product information. For example, those in the tablet market need not hunt around to learn about the product attributes of Apple's tablet. Apple is happy to gather the information and attach it to the recognizable iPad mark, enabling potential purchasers to find the information that interests them (e.g., how much memory does the iPad Air have?). And once the mark is at work, others may chime in with additional information in the form of product reviews and the like. In time, a mark may accumulate meanings having little to do with product attributes. Brands may, for example, evoke a personality independently of any underlying good or service (e.g., the tendency of fans to see themselves as part of a larger community, Red Sox Nation, for example). The process depends, nonetheless, on the mark's ability to communicate a simple signal of source.²⁷

This potential scope of information is both a cost and a benefit to consumers. Trademarks simplify a broad range of meanings; RED SOX, for example, evokes a team, a fandom, and a regional identity. Wearing a branded cap with the team logo may therefore communicate, "the American League baseball team that is based in Boston and plays in Fenway Park," just as it represents "the team for which Ted Williams and Pedro Martinez played" and "the traditional dominant regional sports team of New England." Depending on who wears it and the context, the logo on a cap might also communicate "I'm a Red Sox fan," "I'm from Boston," or even—depending on where I am—"I'm a liberal."

In all these cases, the mark reduces context to simple signals, each of which may have considerably more nuance if spelled out. This can create issues if a mark is not an "empty vessel" but rather brings meaning to the table.²⁸ For example, LOVEE LAMB could not be registered for seat covers because the term suggests a product made from animal skin when in fact it is not.²⁹ Though the seller accurately described the product in advertising, consumers are allowed to rely on the messages carried by the mark, without being expected to hunt for corrective information.³⁰ This simplification also occurs even when a mark is performing a purely source-identifying function. APPLE represents the source of many different kinds of computers in a variety of markets, but the single mark spans them all.

One could spell the meanings out more precisely. The mark is not strictly necessary, but it is helpful. We could imagine a world without enforceable trademark rights in which one sees a soda branded COCA-COLA and then must do the work to learn what precisely is meant by the term. Is it the well-known soft drink or something else? Likewise, those in the market for a tablet computer could gather product information themselves. But time and cognitive capacities are limited.

she liked (or disliked) in the past." (citation omitted) (quoting I J. THOMAS MCCARTHY, MCCARTHY ON TRADEMARKS AND UNFAIR COMPETITION 2-3 (3d ed. 1994)); see also William M. Landes & Richard A. Posner, *Trademark Law: An Economic Perspective*, 30 J.L. & ECON. 265, 269 (1987).

²⁷ Michael Grynberg, *Thick Marks, Thin Marks*, 67 CASE WESTERN RES. L. REV. 13, 20-22 (2016).

²⁸ See generally Michael Grynberg, *A Trademark Defense of the Disparagement Bar*, 126 YALE L.J. F. 178, 183-87 (2016).

²⁹ See generally *In re Budge Mfg. Co.*, 857 F.2d 773 (Fed. Cir. 1988).

³⁰ See *id.* at 775.

Trademark law therefore promotes useful transactions by limiting our need to sort through context for routine purchasing decisions.³¹

But this is necessarily a second-best solution. Using the trademark shortcut comes at the expense of delving more deeply into available information. Our limitations of time and cognition make this eminently rational.³² It nonetheless produces less optimal results than would be available in a world without these limitations; by economizing information we necessarily miss things.

This is a well-known problem in trademark law, reflected by the selling power possessed by many marks. Once a brand is the market leader (or sufficiently advertised and thrust before us), it can take advantage of consumer inertia (fed by cognitive shortcuts like the availability heuristic), making it harder for new entrants to acquire market share.³³ This may be so even when a newcomer offers a superior price/quality balance to the incumbent market leader.³⁴ Likewise, the owner of an established mark may seek to leverage its goodwill in one market to enter

³¹ Nicholas S. Economides, *The Economics of Trademarks*, 78 TRADEMARK REP. 523, 526–27 (1988) (“The economic role of the trademark is to help the consumer identify the unobservable features of the trademarked product. This information is not provided to the consumer in an analytic form, such as an indication of size or a listing of ingredients, but rather in summary form, through a symbol which the consumer identifies with a specific combination of features. Information in analytic form is a complement to, rather than a substitute for, trademarks.”).

³² *Cf. id.* at 524 (“The same symbols can mean different things to different individuals. But, by and large the convention which identifies symbols and words with some minimally defined mental images at a certain point in time for a certain group of people, allows communication and civilization to continue.” (footnote omitted)).

³³ Jeremy N. Sheff, *Biasing Brands*, 32 CARDOZO L. REV. 1245, 1288 (2011). See Amos Tversky & Daniel Kahneman, *Judgment under Uncertainty: Heuristics and Biases*, SCI., Sept. 27, 1974, at 1124, 1130, https://www.socsci.uci.edu/~bskyrms/bio/readings/tversky_k_heuristics_biases.pdf [<https://perma.cc/3KN8-6P9C>], for a discussion of the availability heuristic, which is our tendency to accord greater weight to information that comes immediately to mind. This issue was known before trademark scholars had the language of psychology to describe it. As Ralph Brown wrote decades ago:

With time, the symbol comes to be more than a conduit through which the persuasive power of the advertising is transmitted, and acquires a potency, a “commercial magnetism,” of its own. One of the oldest of advertising techniques, the simple reiteration of the brand name, contributes to this result. Early advertising artists aspired to deface every natural monument with such forgotten symbols as “Sapolio.” Their successors, no longer earthbound, write the bare syllables “Pepsi-Cola” in the sky. If those who crane their necks at the sky-writing are unable to blurt any name but Pepsi-Cola to the soda-clerk, the symbol obviously has commercial value. Even though its continued nurture requires continued outlays, the distillation of past displays and jingles and art exhibits into a word makes that word of great price, quite independently of the vats and alchemy that produce the drink.

Ralph S. Brown, Jr., *Advertising and the Public Interest: Legal Protection of Trade Symbols*, 57 YALE L.J. 1165, 1187–88 (1948) (footnotes omitted).

³⁴ As with the case of generic pain relievers like acetaminophen and naproxen sodium compared to the respective brands of TYLENOL and ALEVE. Sarah Kliff, *Shop like a Pharmacist: Don’t Buy Advil*, VOX (May 10, 2016), <https://www.vox.com/2014/7/25/5936739/shop-like-a-pharmacist-dont-buy-advil> [<https://perma.cc/L6BW-3JXZ>]; see also Julia Belluz, *Stop Wasting Money on Brand-name Drugs*, VOX (Feb. 16, 2016, 9:40 AM), <https://www.vox.com/2016/2/16/11008134/generic-drugs-safe-effective-cheaper> [<https://perma.cc/MGC9-UX3C>] (“The existing body of high-quality evidence suggests that generic drugs consistently meet [the FDA’s equivalence-to-branded-medicine] requirements. So there’s generally little downside to switching to generics. The only difference (in most cases) is that they’re less of a burden on the wallet.”).

another in which the seller lacks comparable expertise. Here too, the availability heuristic might lead consumers to a suboptimal result.³⁵

Another error emanates from the prospect that the trademark carries too much information to be interpreted properly.³⁶ It may be helpful to use a brand as a proxy, but at the end of the day, an APPLE MacBook is not an APPLE MacBook AIR is not an APPLE MacBook Pro is not an APPLE iPhone 7 (or 8 or 9 or X) is not an APPLE iPad, and so on. Strong marks may cause us to overlook distinctions of this sort.

Courts face similar challenges in adjudicating trademark cases. Judges routinely wrestle with the issue of whether to protect not only a mark's source-identifying function, but also its more extended meanings. Because a trademark (or a lookalike) conveys a range of potential meanings, courts sometimes worry about the possibility of harm from non-source messages that might be conveyed by a mark.³⁷ Even when these stories of harm are open to doubt, courts may feel pressure to credit them lest they undermine the overarching structure of trademark law.³⁸

But the AI of the thought experiment could address these concerns. It need not economize on search costs in the same way humans do, for the source of its utility is its superior ability to sift through the context that we seek to avoid. It would know the difference between the many different products sharing, say, the APPLE mark and base recommendations accordingly. Likewise, you may be nervous about trying a new brand when your current choice satisfies well enough, but an AI would be designed for the labor of calculating whether the risk is worth it.

An AI could also account for some of the issues that lead courts to grant strong trademark rights even in the absence of potential diverted sales. Consider *Maker's Mark Distillery, Inc. v. Diageo North America, Inc.*³⁹ *Maker's Mark*, which makes bourbon, sued a tequila producer for using a red wax seal purportedly similar to its own (which is registered as a trademark).⁴⁰ Notwithstanding differences in

³⁵ For example, if COCA-COLA were to release a beer, consumers might transfer their good feelings for the soda to the new product and try it. If the product is poor, they suffer a form of harm inflicted by the mark's selling power. The "error" arises from bounded rationality. It might not have been worth the trouble to thoroughly explore whether there was any reason to think that the soda maker is capable of making a quality beer. The trademark shortcut could blind us to more obvious contextual data, as when a consumer who enjoys COCA-COLA's sweet taste buys COCA-COLA beer without noting that it is marketed as a bitter beer. To be sure, there may be a market check on this kind of conduct. The data is unclear, however, as to whether consumers would actually punish a trademark holder in its home market. See Mark A. Lemley & Mark P. McKenna, *Owning Mark(ets)*, 109 MICH. L. REV. 137, 140-41 (2010) ("[T]he empirical evidence confirms both that third parties can benefit from uses of known marks in markets ancillary to the senior mark owner's and that those third-party uses can impair the senior user's ability to expand its own product lines. Put another way, the evidence suggests that third parties like Black & Decker might benefit from use of, or proximity to, SUM's trademarks, but not that SUM is harmed by such use." (footnotes omitted)).

³⁶ Cf. Brown, *supra* note 33, at 1189 (observing that marks may "be the vehicle of persuasion, either because of extensive repetition and embellishment apart from their use on goods, or because the advertiser has selected and somehow appropriated to his exclusive use a symbol which independently predisposes the customer to buy").

³⁷ See, e.g., *Maker's Mark Distillery, Inc. v. Diageo N. Am., Inc.*, 679 F.3d 410, 419 (6th Cir. 2012); *Anheuser-Busch, Inc. v. Balducci Publ'ns*, 28 F.3d 769, 772-73 (8th Cir. 1994).

³⁸ See Grynberg, *supra* note 27, at 46, 51-52.

³⁹ *Maker's Mark*, 679 F.3d at 410-25.

⁴⁰ *Id.* at 414, 417.

products,⁴¹ price points, and labels, all of which limited the risk of a mistaken purchase due to source confusion, Maker’s Mark prevailed.⁴² The court perceived an intolerable potential for affiliation confusion—the prospect that consumers might see a similar wax seal and muse about the prospect of a connection between the two products—noting that “many consumers are unaware of the affiliations between brands of distilled spirits, and that some companies produce multiple types of distilled spirits.”⁴³

There is much to criticize in the court’s speculations,⁴⁴ but let us assume *arguendo* that: a) consumers might indeed make this leap; *and* b) that the potential misperception of affiliation would be material to a purchasing decision.⁴⁵ Although information about the actual connection, if any, between the two producers is likely publicly available, we do not expect a potential purchaser to look for it while in the aisles of a liquor store. Even if everyone had a smart phone and data plan, web searches take time. We there allow reliance on a simple trademark signal. But sorting out these questions would be trivial for an AI.

An AI, moreover, could transcend the need for this information. To see why, we should ask why affiliation information might be relevant.⁴⁶ Why would a tequila consumer care if there were a connection between a potential purchase and a bourbon producer? After all, shouldn’t the question turn on the tequila’s qualities? One possible answer is that affiliation with a quality bourbon is a proxy for quality. On this logic, if Maker’s Mark makes a quality bourbon, it will protect its goodwill by being careful in entering into affiliation agreements. If so, all things being equal, a tequila affiliated with Maker’s Mark is likely to satisfy Maker’s Mark fans.

But there is typically much better evidence available from other sources—reviews, seller product information, message board discussions, etc. That data is, however, more expensive as a matter of a consumer’s search costs than drawing inferences by simply looking at a mark. Maybe it is easier to make assumptions based on bottle appearance. An AI would alter the equation. Because it could accumulate data at a low cost, it would have no need to rely on the trademark’s ability to convey affiliation information. In short, it would have immediate access to: (a) accurate affiliation information if relevant; and (b) superior information that might make affiliation data irrelevant.⁴⁷

⁴¹ Although tequila is a spirit, it is not distilled from a grain like bourbon, which must have a majority-corn base. 27 C.F.R. § 5.22(b)(1)(i), (g); *see also* *Maker’s Mark*, 679 F.3d at 415, 423.

⁴² *Maker’s Mark*, 679 F.3d at 414, 423–25.

⁴³ *Id.* at 422. For that reason, the court discounted the presence of the defendant’s entirely distinct house mark. *Id.* (agreeing with the lower court that “the presence of a house mark . . . is more significant in a palming off case than in an association case”).

⁴⁴ Grynberg, *supra* note 27, at 41–45.

⁴⁵ Though there is no materiality requirement in infringement litigation, materiality in this case would go a long way to justifying the outcome as a matter of policy.

⁴⁶ To be sure, I believe the relevance to be minimal outside of direct sponsorship settings—that is situations in which the markholder stands behind the quality of a third party’s goods—but I am assuming *arguendo* that it matters.

⁴⁷ This example shows a potential pitfall of a less advanced AI that might pay too much attention to trademarks rather than sorting through the more directly relevant context and specific information for

B. *What's left?*

The thought experiment thus highlights some of the tradeoffs underlying today's trademark system. Trademarks help organize information into simple signals that promote efficiency at the expense of nuance and context. The balance, coupled with the law that protects it, is likely helpful to consumers when calibrated to a given level of human cognition. But the proper scope of trademark law may change if new tools enhance our effective cognition. Change the assumptions enough, and much of what makes trademark attractive—information simplification—looks less like the solution to a problem and more like a problem to be solved.

What then happens to trademark law in a world of extremely sophisticated AI? What would be left for it to do?

i. Maintaining distributive fairness

Of course, an AI future may not (and if past is prologue, will not) be equitably distributed. The demands of distributive justice may therefore create a potential foothold for traditional trademark law to persist. Perhaps the hypothesized AIs will not be available to all or maybe a significant amount of commerce will persist in realms that the AI cannot easily operate (e.g., face-to-face interactions).⁴⁸ This could be seen as another issue of intermediate technology. Instead of addressing an AI that is advanced but imperfect, this question is of an AI that is perfect enough, but unevenly distributed.

Though distributive concerns may indeed preserve room for trademark law to operate, the amount of room may be less than might initially appear. First, if we are concerned with life online (or the future equivalent) then the wealth gap issue might at least be mitigated by the relative low cost of digital technology compared to physical goods (assuming that difference persists in the future).⁴⁹ And if life *offline* is the issue, trademark law could be calibrated to operate with greater scope offline than on.⁵⁰ This calibration could also attend to the interests of those for whom the act of shopping is a pleasure unto itself (though perhaps these are consumers for whom sifting through context is part of the fun).

Second, even if top-of-the-line AIs are imperfectly distributed, lesser and more available technologies may be good enough to justify weakening trademark law. This assumes that the weakening translates into an improved marketplace. From a welfare perspective any trademark costs to consumers must be weighed against potential

which the marks are an imprecise stand-in.

⁴⁸ But who knows what developments will arise with augmented reality technology.

⁴⁹ Though perhaps this opens the door to battles as sophisticated AIs seek to deceive other sophisticated AIs. Maybe this leaves room for trademark to operate, but the requisite judgments may be better suited to false advertising law. *See infra* Section II.B.3.

⁵⁰ The early experience with trademarks on the internet was, however, the opposite. *See generally* Eric Goldman, *Brand Spillovers*, 22 HARV. J.L. & TECH. 381, 383 (2009) (“[U]nluckily retailers’ use of brand spillovers, online brand spillover activities have been repeatedly attacked in courts and legislatures.”).

benefits in the actual marketplace (e.g., by lowering licensing costs and, therefore, prices).⁵¹

ii. Preserving a residual core

Another possibility is trademark law will still be needed to ensure that consumers get what their AIs order. But this residual core of trademark law, if necessary, is far narrower than today’s doctrine. The kind of meaning that would need protection—designation of source at the point of sale/delivery—is a faint shadow of what trademark law protects today. In the AI world, there is no need for brand personality, dilution protection, affiliation or sponsorship claims, or the like. Nor is there a need—outside of the prestige goods context discussed below⁵²—for attractive or memorable marks. To a computer, ZL3XC!7K4BV functions just as well as APPLE. But because the AI is able to find quality goods (however defined), the seller retains an incentive to invest in quality (and, indeed, might have more resources to do so if freed from a need to invest in the now irrelevant attribute of trademark attractiveness).⁵³ Or trademark law could be restricted to a smaller signifier. So anyone could brand their computer APPLE, but only one company could use the ®, or some like symbol, when it comes time to ship.

Nor is it entirely clear that an AI could not distinguish authentic from non-authentic goods even without a trademark residual. Imagine seeing luxury branded merchandise for sale on a street corner. No matter how authentic it looks, the context likely alerts you to the strong possibility that the products are counterfeits. The AI would likely have access to considerably more distinguishing context and the ability to sort it.

iii. Everything is false advertising

Perhaps there will be more work, relatively speaking, for the Lanham Act’s false advertising cause of action.⁵⁴ Presumably, advertising would continue to shape the consumer preferences that constitute an input to the AI’s work, leaving room for false advertising doctrine to operate.

One question is the relative balance of trademark and false advertising law in efforts by sellers to “game” imperfect AIs, perhaps by using powerful computers of

⁵¹ See *infra* Section II.C.

⁵² See *infra* Section II.B.4. Mark attractiveness might be an issue to the extent scarcity or cachet is the relevant “product,” but those attributes can be manufactured in other ways.

⁵³ To be sure, however, one consequence might be that marks may lose the placebo effect that comes from mark strength. See, e.g., Kate Faasse et al., *Impact of Brand or Generic Labeling on Medication Effectiveness and Side Effects*, 35 HEALTH PSYCH. 187, 187 (2016). While this is often seen as reflective of manipulation by marketing, a recent article sees the placebo effect of trademarks as potentially salutary. Jake Linford, *Placebo Marks*, 47 PEPP. L. REV. 45, 112–13 (2019) (“These studies suggest that in the market for high-performance goods, consumers may derive value from the branding myths they are sold. The Nike brand may work like Dumbo’s feather in the famous Disney film—it may not matter why consumers believe they can fly, so long as they believe it.”). In any case, the authority of the AI could well imbue its choices with a placebo effect of its own.

⁵⁴ See 15 U.S.C. § 1125(a) (2018).

their own. The differing foci of the two doctrines suggest that false advertising will have a greater role to play in an AI world—at least if the concern is with deception. Trademark law is oriented around simple signals that embody broader meanings. False advertising law, by contrast, looks to the message as a whole.⁵⁵ It is therefore more context sensitive.⁵⁶

This orientation would better serve efforts to police activities that mislead AIs. Because the hypothesized AIs focus on context and not trademarks alone, mere trademark use is comparatively unlikely to mislead an AI of sufficient sophistication. We might, however, posit circumstances in which surrounding context is fabricated in a way designed to game the AIs. False advertising law, with its demand that actionable communications be both perceived and material, is better positioned to respond to acts that might result in mistaken purchases.

iv. Everything is unfair competition and the problem of prestige goods

Or trademark law might discard any pretense of focusing on consumer protection and return to the unfair competition tradition.⁵⁷ Judges might adjudicate what is or is not acceptable behavior in the commercial marketplace without necessarily focusing on consumer confusion. This could be the mechanism by which they continue to police competition in status goods or promotional merchandise, with underlying views of what is “sporting” in the marketplace, substituting for today’s strained stories about possible consumer confusion.⁵⁸ But without confusion, these moral judgments will need a coherent theory if the resulting law is to be transparent and predictable.⁵⁹

⁵⁵ See *Church & Dwight Co. v. SPD Swiss Precision Diagnostics, GmbH*, 843 F.3d 48, 65 (2d Cir. 2016) (“To prevail on a Lanham Act false advertising claim, a plaintiff must establish that the challenged message is [] either literally or impliedly false . . .”).

⁵⁶ Moreover, the doctrine makes distinctions between explicit and implied falsehoods, requiring plaintiffs to establish that consumers actually perceive the implied falsehood. *E.g.*, *Time Warner Cable, Inc. v. DIRECTV, Inc.*, 497 F.3d 144, 153 (2d Cir. 2007). The doctrine further considers context by requiring materiality and excluding from liability statements that are mere puffery. *E.g.*, *Pizza Hut, Inc. v. Papa John’s Int’l, Inc.*, 227 F.3d 489, 495–96 (5th Cir. 2000).

⁵⁷ Passage of the Lanham Act in 1946 unified federal trademark protection in one body of law. See 1 J. THOMAS MCCARTHY, MCCARTHY ON TRADEMARKS AND UNFAIR COMPETITION § 4:3 (5th ed. 2019). Federal statutory protection was once restricted to “technical trademarks,” which would be called arbitrary and fanciful marks today. *Id.* The law of unfair competition covered what were then known as trade names but are now treated as trademarks. *Id.* Both types of marks are now protected by the Lanham Act. *Id.*

⁵⁸ As when the Ninth Circuit used trademark law to prevent the use of popular trademarks for automobiles as raw material for complementary goods such as keychains or license-plate frames. *Au-Tomotive Gold, Inc. v. Volkswagen of Am., Inc.*, 457 F.3d 1062, 1064 (9th Cir. 2006). Though the court undertook its analysis under the multifactor test, *id.* at 1078, the opinion seems more shaped by the reaction that the challenged use was “nothing more than naked appropriation of the marks,” *id.* at 1064.

⁵⁹ It is worth noting, however, that trademark law may be making room for explorations of this sort. The *Lexmark International, Inc. v. Static Control Components, Inc.* decision may breathe new power into the Lanham Act’s reference to the prevention of “unfair competition” as a statutory purpose. See *Lexmark Int’l, Inc. v. Static Control Components, Inc.*, 572 U.S. 118, 131–32, 136 (2014). Recent precedent from the Fourth Circuit, moreover, suggests that courts may be interested in exploring a law of unfair competition that is broader than trademark alone. See *Belmora LLC v. Bayer Consumer Care AG*, 819 F.3d 697, 701, 706–08 (4th Cir. 2016).

Judges already use trademark law to regulate competition in a variety of areas in which source confusion is unlikely. Courts have used the Lanham Act to give trademark holders control of the promotional goods/merchandising market (e.g., preventing third parties from selling baseball caps with professional team logos) and the market for prestige goods (e.g., using post-sale confusion theories to prevent the selling of knockoff products even when the buyer is aware that the product does not come from the trademark holder).⁶⁰

These extensions of trademark law are controversial for a number of reasons, and they rest doctrinally on confusion theories that generally fall outside of the source confusion model.⁶¹ In the hypothesized AI world, in which confusion is eliminated by automatic consideration of context, these theories would be less viable. Perhaps a result will be the lower prices for consumers that come with competition.

Alternatively, today’s prohibitions would remain, but they might be liberated from their strained confusion stories. Defenders of post-sale confusion and merchandising rights theories generally do not rely on confusion rationales. Rather, they focus on considerations of promoting incentives,⁶² deterring free riding,⁶³ and policing morality.⁶⁴ These rationales are simply unnecessary to a trademark law grounded in preventing passing off and source confusion, and they have always been an uneasy fit with trademark doctrine as a whole. But in a world in which source confusion is harder to come by, they may become the new core of trademark law. If the link to source confusion is severed, perhaps the resulting law of luxury and promotional markets could be founded on a non-confusion theory that offers greater clarity to future judges of what they are trying to accomplish.

C. *What’s gained?*

To say that an AI-enabled world of consumer consumption does not need trademark law is not the same as saying we would be better off without it. What then might be gained by scaling back trademark law as we know it today? Here, too, much depends on how one imagines the state of intermediate technologies and their distribution.

i. Legalizing competition

As suggested in the last subpart, a major potential benefit would be the legalization of products and services threatened or prohibited by current trademark law and lower prices in prestige or promotional merchandise.⁶⁵ As noted above, however, the intuition that these markets “belong” to trademark holders is powerful

⁶⁰ E.g., *Hermès Int’l v. Lederer de Paris Fifth Ave., Inc.*, 219 F.3d 104 (2d Cir. 2000); *Boston Prof’l Hockey Ass’n v. Dallas Cap & Emblem Mfg.*, 510 F.2d 1004 (5th Cir. 1975).

⁶¹ See, e.g., *Hermès*, 219 F.3d at 108; *Boston Prof’l Hockey*, 510 F.2d at 1012.

⁶² See, e.g., *Boston Prof’l Hockey*, 510 F.2d at 1011.

⁶³ See, e.g., *Au-Tomotive Gold*, 457 F.3d at 1064, 1067.

⁶⁴ See, e.g., *Hermès*, 219 F.3d at 108–09.

⁶⁵ See *supra* Section II.B.4.

and likely to generate another cause of action to replace the one that is lost.⁶⁶ Even this may be a gain, as a new cause of action would be liberated from incoherent theories of consumer confusion.

ii. Opening communications channels

Another potential benefit of weakening trademark in an AI world is that sellers would have the benefit of enhanced communication channels to consumers. A maker of acetaminophen could, for example, simply brand itself as TYLENOL, leaving it to consumer avatars to sort out which is the “authentic” brand for purchasers who actually care about the difference. Because the brand name conveys useful product information, weakened trademark law would enable an easy way of communicating a product attribute (as TYLENOL conveys non-brand information that “pain reliever” does not).

This point should not be oversold. A world in which AI confusion is impossible is one in which the AI would be alert to the prospect of generic brands. In today’s market, for example, a price-sensitive consumer who is indifferent to whether her acetaminophen is TYLENOL may or may not know of the availability of generic alternatives. Our hypothesized AI would not suffer this difficulty and would simply order the cheaper generic equivalent for its consumer (assuming that it would satisfy the consumer’s preference). An AI would not be put off by the comparative clunkiness of the term acetaminophen.⁶⁷

Another possible answer would be to focus not on the consumer-side AI but rather the seller’s advertising costs. So a seller may wish to use a trademarked term as the most efficient way to communicate the non-trademark information that the term embodies (in the drug example, “my product is like TYLENOL”). If forced to incur the costs of advertising without the trademarked term, the seller may have to raise prices. But again, if the AIs are good enough, the costs of marketing to them should likewise be limited.⁶⁸

Perhaps a more promising answer is to focus on trademark litigation’s potential as a weapon between market competitors. Even if sellers do not “need” plaintiff marks, many legitimate activities may nonetheless provoke trademark claims. A seller may, for example, design a product that looks similar to a claimed trade dress

⁶⁶ See *supra* Section II.B.4.

⁶⁷ But there is an effect in a world of intermediate AIs, which may be gamed on the one hand, or fail to seize opportunities on the other. Here, confusion costs would have to be weighed against benefits. See generally Michael Grynberg, *Trademark Litigation as Consumer Conflict*, 83 N.Y.U. L. REV. 60 (2008).

⁶⁸ We might also see the issue as one of speech protection. A “counterfeit” mark might provide relevant information to the AI. Suppose you are in the market for a cheap watch, and you see three \$15 options: A well-known discount brand, an unknown cheap brand, and an obviously fake ROLEX. The counterfeit use of ROLEX (which does not deceive you) conveys information about the product that bears it.

If the AI is incapable of confusion, then stopping “counterfeits” becomes another way of saying, “don’t tell the machine this,” rather than a means of preventing harmful deception. This raises First Amendment concerns. Cf. *United States v. Alvarez*, 567 U.S. 709, 734 (2012) (Breyer, J., concurring) (noting that statutes and doctrines that prohibit false statements “limit the scope of their application, sometimes by requiring proof of specific harm to identifiable victims; sometimes by specifying that the lies be made in contexts in which a tangible harm to others is especially likely to occur; and sometimes by limiting the prohibited lies to those that are particularly likely to produce harm”).

for functional reasons. Or a seller may use a trademarked term in a descriptive way. These legitimate activities create openings for trademark-based attacks.

Today, such cases require courts to sort through difficult questions, e.g., whether product designs have sufficient secondary meaning for trademark protection, whether designs are functional, whether consumers are likely to be confused by similar competitor designs, and so on. To the extent the underlying claims are fundamentally about preserving market share—rather than consumer protection—limiting them would promote competition.⁶⁹

iii. Protecting free expression

A final benefit of ending trademark law as we know it is eliminating infringement litigation (or its threat) as a tool of censorship of expressive trademark uses. The unauthorized use of trademark for expressive purposes may benefit the public by, for example, enriching culture or providing an avenue for commentary. For many of these uses—unlike the traditional passing off setting—the junior user is unable to adopt a non-confusing distinctive mark. Using an already-known trademark is part of the message.

Current trademark doctrine mediates the potential interference with free expression in a variety of ways.⁷⁰ Though generally effective when push comes to shove,⁷¹ trademark law’s safeguards for free expression may not be enough to blunt the *in terrorem* effect of a cease and desist letter for those with limited access to legal advice. And those without a taste for litigation may simply trim their sails in the face of a plausible sounding trademark claim.⁷² If courts can trust the AI, the technology would eliminate this tactic.

D. The “death of trademark” in a bespoke world

As the saying goes, predictions are hard, especially about the future,⁷³ but the thought experiment helps illuminate the functions of trademarks today. Trademarks exist in a world of limited cognition. They, and the law that protects them, are one way of addressing marketplace information problems. But they are not mandatory. Our use of trademarks is the product of a particular context. If and when our abilities improve, trademarks may become superfluous. At the least, the current equilibrium that justifies a particular scope for trademark law may reset at another level.

⁶⁹ See *Wal-Mart Stores, Inc. v. Samara Bros.*, 529 U.S. 205, 214 (2000) (noting the importance of providing an avenue “for summary disposition of an anticompetitive strike suit”).

⁷⁰ See, e.g., *New Kids on the Block v. News Am. Publ’g, Inc.*, 971 F.2d 302, 308 (9th Cir. 1992) (protecting nominative uses from liability); *Rogers v. Grimaldi*, 875 F.2d 994, 998–1000 (2d Cir. 1989) (protecting artistically relevant uses from liability).

⁷¹ *But see* *Gordon v. Drape Creative, Inc.*, 909 F.3d 257, 271–72 (9th Cir. 2018) (remanding trademark claim for further proceedings notwithstanding existence of an artistically relevant use by the defendant).

⁷² *Cf. Wal-Mart*, 529 U.S. at 213–14; see generally James Gibson, *Risk Aversion and Rights Accretion in Intellectual Property Law*, 116 *YALE L.J.* 882, 913 (2007).

⁷³ *It’s Difficult to Make Predictions, Especially About the Future*, QUOTE INVESTIGATOR (Oct. 20, 2013), <https://quoteinvestigator.com/2013/10/20/no-predict/> [<https://perma.cc/ZF8V-JKVJ>].

It turns out the current trademark equilibrium is already under pressure from artificial intelligence. Although the AI of the thought experiment is not at hand, its forerunners can be seen in the marketplace, and they already affect trademark law, as discussed in Part IV below.⁷⁴

But the thought experiment highlights another tension within modern trademark doctrine. The hypothetical AI is a digital creature at home on the internet. It acts not by relying on any particular source of information, but rather by filtering the mass of data available online. As the next Part explains, that model is at odds with how information—including trademark information—was managed prior to the advent of the internet.

III. KNOWLEDGE ONLINE AND TRADEMARKS

In our hypothetical, the AI avatar is a filter. It sorts and categorizes the myriad of resources available online and distills a message on behalf of its beneficiary consumer. This is a familiar way to gather information online. If interested in exploring a subject, we might rely on a filter (like Google) to narrow the range of possibilities before us. But the filtered out possibilities remain available to those using alternative screens or search terms.

Information gathering was different in a world of comparative information scarcity. No newspaper could contain all the stories of the day; no library all the books; no television channel all the shows. Choices had to be made, and with the necessity of choice came the requirement of authorities—the editors, librarians, and station managers—to curate what we could see.⁷⁵

The powerful forces unleashed by the internet changed that. In many ways, however, trademark law functions in a way akin to the old scarcity model, creating challenges for its accommodation to life online.

A. Knowledge Online and Trademarks

How do people come to rely on the opinions of others and find the information on which they rely? For many, the abundance of information online changes the answer to this question by altering the way knowledge is aggregated and curated. The role of gatekeepers used to be larger.⁷⁶ We once consumed information from a

⁷⁴ See *infra* Part IV.

⁷⁵ David Weinberger describes how the physical limits of communications media shaped the institutions that used them:

Traditional knowledge is what you get when paper is its medium. There is nothing mystical about this. For example, if your medium doesn't easily allow you to correct mistakes, knowledge will tend to be carefully vetted. If it's expensive to publish, then you will create mechanisms that winnow out contenders. If you're publishing on paper, you will create centralized locations where you amass books. . . . Traditional knowledge has been an accident of paper.

WEINBERGER, *supra* note 2, at 45.

⁷⁶ See *id.* at 3–4.

set of relatively constrained options, choosing what to view only after others had, by necessity, made a great many choices. As David Weinberger explains:

Our most basic strategy for understanding a world that far outruns our brain’s capacity has been to filter, winnow, and otherwise reduce it to something more manageable. We’ve managed the fire hose by reducing the flow. We’ve done this through an elaborate system of editorial filters that have prevented most of what’s written from being published, through an elaborate system of curatorial filters that has kept most of what’s been published from being shelved in our local libraries and bookstores, and through an elaborate system of professional filters that have kept many of us from being responsible for knowing most of what’s made it through the other filters. Knowledge has been about reducing what we need to know.⁷⁷

This is largely a byproduct of scarcity. There are only so many library shelves, pages in a newspaper, broadcast television channels, etc. Publishing houses could print only so many books. Radio and television broadcast networks were expensive to create, and the transmission spectrum was managed as a scarce resource.⁷⁸ And so on.

Enter the experts. Authorities (libraries, universities, newspapers, etc.) acted as curators of knowledge, effectively determining the bounds of the conventional wisdom.⁷⁹ One need not ascribe censorial motives to this arrangement. The existence of gatekeepers, of experts curating “correct” answers, was inherent to the prevailing technology—a world of limited shelf space requires choices.⁸⁰

On the internet, of course, things are different. Scarcity of supply is not a problem, as for all practical purposes the shelf space is infinite. The problem is finding one’s preferred signal amid the surrounding noise.⁸¹ We therefore need a different class of experts. Instead of depending on librarians and editors who select from the submissions of information providers and leave the rest unavailable, we rely on filters, algorithms, and social networks to predict what we want to see and leave the remaining mass intact and undifferentiated.⁸² But the “wrong” answers remain available for those who see them as correct. The internet in turn connects adherents to these alternative premises and forms communities that resist the ability of the old gatekeepers to make effective pronouncements on contested issues.⁸³

⁷⁷ *Id.* at 3–4.

⁷⁸ Gregory Staple & Kevin Werbach, *The End of Spectrum Scarcity*, SPECTRUM (Mar. 1, 2004, 15:16 GMT), <https://spectrum.ieee.org/telecom/wireless/the-end-of-spectrum-scarcity> [<https://perma.cc/D6BD-KQB8>].

⁷⁹ See WEINBERGER, *supra* note 2, at 3–4. This is a comparative point, as there has always been dissenting opinions and narratives.

⁸⁰ *See id.* at 5.

⁸¹ *See id.* at 45 (“The Internet is what you get when everyone is a curator and everything is linked.”). In some respects, this is nothing new, as “[t]here was always too much to know, but now that fact is thrown in our faces at every turn. Now we *know* there’s too much for us to know.” *Id.* at 11.

⁸² *See id.* at 11 (“Filters no longer filter out. They filter *forward*, bringing their results to the front. What doesn’t make it through a filter is still visible and available in the background.”).

⁸³ *See id.* at 45 (“The Internet simply doesn’t have what it takes to create a body of knowledge: No editors and curators who get to decide what is in or out. No agreed-upon walls to let us know that knowledge begins here, while outside uncertainty reigns—at least none that everyone accepts.”).

Democratizing knowledge in this way has not been an unalloyed good, as seen in the 2016 presidential election campaign.⁸⁴ Multiple parallel knowledge communities may develop world views from their own selected body of facts, selected in a filtering process by members of a network.⁸⁵ The all-too-familiar problem of ideological filter bubbles is one result. Partisan Democrats and Republicans perceive reality in very different ways as do those on opposite sides of debates over gun control or climate change. This is also true of smaller groups. Many communities organize and define themselves around premises that are fundamentally irreconcilable with those held by mainstream society. Some of these may seem harmless enough to just be fodder for ridicule (e.g., flat earthers).⁸⁶ Others carry higher emotional stakes (e.g., “9/11 ‘truthers’” and “birthers”)⁸⁷ and cause direct, real-world harm (e.g., disease outbreaks traceable to anti-vaccine misinformation).⁸⁸

Regardless of the consequences of online communities holding incompatible visions of truth, that reality seems here to stay.⁸⁹ No amount of argument, no matter the consensus in society at large, is going to convince the birther that President Obama was born in the United States. More data isn’t going to convince the anti-vaccine crusaders of the dangers of their cause no matter the costs to the rest of us.⁹⁰

But this is an article about trademarks. How does trademark law fit into the evolution described above? I have two points. First, trademark information has traditionally been managed in a manner consistent with the scarcity model. Second, the conflicting internet model of information abundance places pressure upon trademark law. Trademark law seeks to manage the shelf space, so to speak, by constraining meanings associated with a mark. The internet model would allow the proliferation of meanings, relying on filters to distinguish “authorized” views (i.e., those approved by trademark owners) from conflicting information, which remains available for those who want it. In this way, the internet model foreshadows the world of the AI hypothesized in the previous parts. The closer technology gets to making it a reality, the weaker trademark law may become.

⁸⁴ See generally, e.g., DAVID E. SANGER, *THE PERFECT WEAPON* 171–93 (1st ed. 2018).

⁸⁵ See WEINBERGER, *supra* note 2, at 10–12.

⁸⁶ See John Timmer, *Why Does Flat Earth Belief Still Exist?*, *ARS TECHNICA* (Dec. 28, 2018, 10:30 AM), <https://arstechnica.com/science/2018/12/why-does-flat-earth-belief-still-exist/> [<https://perma.cc/ZS56-7WYG>].

⁸⁷ See, e.g., Chris Bell, *The People Who Think 9/11 May Have Been an ‘Inside Job,’* *BBC: NEWS* (Feb. 1, 2018), <https://www.bbc.com/news/blogs-trending-42195513> [<https://perma.cc/FAS3-XCSB>]; Lily Rothman, *This is How the Whole Birther Thing Actually Started*, *TIME* (Sept. 16, 2016), <https://time.com/4496792/birther-rumor-started/> [<https://perma.cc/VAT2-KMAR>].

⁸⁸ See Elizabeth Cohen & John Bonifield, *Her Son Died. And then Anti-vaxers Attacked Her*, *CNN: HEALTH* (Mar. 21, 2019, 2:47 PM), <https://www.cnn.com/2019/03/19/health/anti-vax-harassment-eprise/index.html> [<https://perma.cc/QMF8-MXKR>]; cf. WEINBERGER, *supra* note 2, at 151 (“Science is not going to be able to reassert its old-style authority because it has lost the medium that enabled it to flourish: a one-way channel in which there were those who spoke and those who listened.”).

⁸⁹ See WEINBERGER, *supra* note 2, at xiii (“[W]e are in a crisis of knowledge at the same time that we are in an epochal exaltation of knowledge. We fear for the institutions on which we have relied for trustworthy knowledge, but there’s also a joy we can feel pulsing through our culture. It comes from a different place. It comes from *the networking of knowledge*.”).

⁹⁰ *Id.* at 181 (“There is nothing you can say to convince some people. The old Enlightenment ideal was far more plausible when what we saw of the nattering world came through filters that hid the vast, disagreeable bulk of disagreement.”).

This reflects a tension in trademark jurisprudence that began in the late 1990s. Some judges initially doubled down on the scarcity model, expanding the authority of trademark holders.⁹¹ As discussed below, however, the logic of the internet information model of abundance has proven too powerful for trademark doctrine to ignore.

B. Trademarks and the stability of meaning

To return to trademark basics, trademarks (and service marks) identify and distinguish goods (and services) in the marketplace and serve consumers by simplifying information.⁹² This function requires relative stability of meaning for the trademark itself. COCA-COLA does not work as a source identifier if the term delineates competing brands of soda. Stated another way, the “space” in the term for source-identifying meanings is limited if a trademark is to perform an information-economizing function.

In the brick-and-mortar world this traditionally meant that someone must define the mark. Trademark law therefore gives mark owners a favored position in curating trademark information by letting them police the use of confusingly similar marks.⁹³ This protects not only source information but also the marks’ ability to embody seller goodwill with consumers.⁹⁴ By being the only ones authorized to designate what the mark’s source meaning represents, sellers manage the goodwill attached to the mark by controlling product quality, planning advertising campaigns, interacting with customers, and the like.

Trademarks are also partially protected from accumulating multiple source meanings as a result of markholder conduct.⁹⁵ Trademarks lose protection when their owners license or assign them indiscriminately or fail to use them for extended periods of time.⁹⁶ So trademark law not only encourages active curation of meaning, it may be said to require it.⁹⁷

⁹¹ See, e.g., *Allard Enters. v. Advanced Programming Res., Inc.*, 249 F.3d 564, 575 (6th Cir. 2001) (vacating district court injunction that would have precluded a concurrent junior user from any internet use of contested mark).

⁹² See *supra* Section II.A.

⁹³ See 1 MCCARTHY, *supra* note 57, at § 2:10

⁹⁴ See *id.* § 2:17 (defining goodwill as “that which makes tomorrow’s business more than an accident. It is the reasonable expectation of future patronage based on past satisfactory dealings” (quoting EDWARD S. ROGERS, *GOOD WILL, TRADEMARKS AND UNFAIR TRADING* 13 (1914))).

⁹⁵ See *id.* § 3:12.

⁹⁶ 15 U.S.C. § 1127 (2018) (providing that a mark is deemed abandoned “[w]hen its use has been discontinued with intent not to resume such use. . . . Nonuse for 3 consecutive years shall be prima facie evidence of abandonment” and when “any course of conduct of the owner” causes the mark “to lose its significance as a mark”).

⁹⁷ Many traditional limits to trademark rights may also be understood as supporting stability of trademark meaning. See Grynberg, *supra* note 28, at 183–87. Marks deemed unlikely to perform a source-identifying function are harder to protect. That is, if a term is unlikely to maintain a discrete meaning in the marketplace, it is less likely to receive protection. Generic terms that identify product categories may never be trademarks while descriptive terms require secondary meaning, that is, consumers must have come to associate the term with a particular product source. *Id.* at 185. Even then, competitors may continue to use the word in its original, non-trademark sense. *Id.* at 187.

As discussed above, modern trademark law goes beyond the protection of source meanings.⁹⁸ This leads to the question of whether protection of non-source meanings is necessary to secure the positive externalities that come with the creation of a source-identifying mark.⁹⁹ For many judges, however, this leap is plausible. Trademark rights therefore now extend to uses that might cause consumer confusion with respect to a markholder's approval or sponsorship as well as activities that might dilute a famous trademark.¹⁰⁰ While there is much to criticize in this expansion, it may be described as upholding the one mark/one meaning view of trademarks.¹⁰¹

But though the scarcity model gives trademark holders ample power, it traditionally contained built-in limitations. Just as a newspaper editor selects articles only for her paper, so too were markholders limited to the geographic markets that had defined the marks.¹⁰² There is no meaning to stabilize where no one has heard of the mark. To hold otherwise, and force junior remote users to abandon their marks, would be to treat the senior user's trademark rights as simple property interests.

Accordingly, at common law, courts limited the scope of protection to the mark's area of use and reputation.¹⁰³ Many early cases took the principle of market definition quite far, limiting infringement actions to cases of direct competition between the trademark holder and the defendant. Today, of course, trademark rights reach beyond the mark's immediate market, but proximity still matters in assessing the likelihood of consumer confusion in an infringement suit,¹⁰⁴ and geographic scope of use still matters absent a registration (which confers nationwide priority).¹⁰⁵

The shift to a system granting nationwide priority for registered marks arguably made trademarks more like a property right, but it reflected a changing view of the national economy. The drafters of the Lanham Act understood this change to be

Likewise, many defensive doctrines that shield potential trademark defendants from liability are designed to insulate a trademark's meaning from distortion. The first-sale (or exhaustion) doctrine allows purchasers of genuine goods to dispose of them as they wish, including by resale. 4 J. THOMAS MCCARTHY, MCCARTHY ON TRADEMARKS AND UNFAIR COMPETITION § 25:41 (5th ed. 2019). In such cases, the trademark on the good retains its meaning—the holder of the mark remains the original source, and the resale does not distort that information. The caveat to the doctrine comes when the reseller undermines that expectation by altering or reconditioning the product. *Id.* Beyond a certain point, courts will declare that the retained trademark no longer designates source in a meaningful way, making its continued use deceptive. *Id.*

The various judicial approaches to protecting nominative trademark uses from liability are to the same effect. A nominative use is one that uses the trademark to refer to its owner, as in comparative advertising (e.g., PEPSI is better than COKE). *Id.* § 23:11. To avoid liability, such uses must not distort the mark's original meaning by making misrepresentations about its owner. So, for example, the Ninth Circuit, which formalizes protection of nominative uses in a "nominative fair use" test, asks whether the defendant "falsely suggested he was sponsored or endorsed by the trademark holder." *Adobe Sys. Inc. v. Christenson*, 809 F.3d 1071, 1081 (9th Cir. 2015).

⁹⁸ See *supra* Section II.A.

⁹⁹ Grynberg, *supra* note 27, at 64–65.

¹⁰⁰ See 15 U.S.C. § 1125(a), (c) (2018).

¹⁰¹ For example, Judge Posner has sought to justify dilution theory by arguing that even non-confusing uses of a mark might interfere with consumer associations with the original mark. See *Ty Inc. v. Perryman*, 306 F.3d 509, 511 (7th Cir. 2002).

¹⁰² *United Drug Co. v. Theodore Rectanus Co.*, 248 U.S. 90, 98 (1918).

¹⁰³ See, e.g., *id.* at 101.

¹⁰⁴ E.g., *Polaroid Corp. v. Polarad Elecs. Corp.*, 287 F.2d 492, 495 (2d Cir. 1961).

¹⁰⁵ See 15 U.S.C. § 1057(c) (2018).

necessary in light of the rise of nationwide markets.¹⁰⁶ Likewise, greater interstate travel made it more likely that consumers would encounter familiar marks outside of their home market.¹⁰⁷ Peaceful, remote coexistence between markholders seemed at risk, and the Lanham Act’s registration system created a mechanism to choose between remote claimants.¹⁰⁸ This logic has, however, troubling implications when applied to the internet.

C. *Stability of meaning online: Two views*

How does the internet affect a trademark’s information function? We might tell two stories—one justifying expanding trademark rights and one contracting them. Both find support in the case law.

i. The internet and the collapse of market barriers

I. *The story*

The internet facilitates the trend to more powerful marks by further breaking down barriers between markets. Just as the rise of national markets supported the Lanham Act’s establishment of nationwide priority for registered marks, the internet likewise combines previously discrete markets. On this logic, the internet makes stronger trademark rights more likely in four ways.

First, being online reduces the importance of spatial borders, uniting remote buyers and sellers irrespective of geography. This facilitates skepticism of a strong territoriality principle, weakening a key limitation to trademark rights. And indeed, some judges have argued that territoriality-based doctrines should have less room to operate in the modern age.¹⁰⁹

¹⁰⁶ See S. REP. NO. 79-1333, at 5 (1946) (“However, trade is no longer local, but is national. Marks used in interstate commerce are properly the subject of Federal regulation. It would seem as if national legislation along national lines securing to the owners of trade-marks in interstate commerce definite rights should be enacted and should be enacted now.”).

¹⁰⁷ See *id.*

¹⁰⁸ See 15 U.S.C. § 1057(c) (2018) (“Contingent on the registration of a mark on the principal register provided by this chapter, the filing of the application to register such mark shall constitute constructive use of the mark, conferring a right of priority, *nationwide in effect*, on or in connection with the goods or services specified in the registration against any other person except for a person [with prior use or registration rights].” (emphasis added)).

¹⁰⁹ For example, the *Dawn Donut* rule provides that even when a registered markholder has priority over a remote junior user, no infringement remedy is possible until there is an actual likelihood of confusion in the market. *Dawn Donut Co. v. Hart’s Food Stores, Inc.*, 267 F.2d 358, 364 (2d Cir. 1959). This requires the plaintiff to commence or be likely to commence activity in the junior user’s market; until then, the junior user may be on borrowed time, but may still operate. See *id.* at 364. This logic has been questioned in light of the growing ability for a mark to establish a reputation in an area absent actual sales. See *Circuit City Stores, Inc. v. CarMax, Inc.*, 165 F.3d 1047, 1057 (6th Cir. 1999) (Jones, J., concurring) (“The *Dawn Donut* Rule was enunciated in 1959. Entering the new millennium, our society is far more mobile than it was four decades ago. For this reason, and given that recent technological innovations such as the Internet are increasingly deconstructing geographical barriers for marketing purposes, it appears to me that a re-examination of precedents would be timely to determine whether the *Dawn Donut* Rule has

Second, the internet weakens conceptual barriers between markets, making it easier for buyers to move between proximate, but distinct, markets. One wanting an economy car is unlikely to wander into a MERCEDES showroom. But one searching for, say, a KIA online may well encounter mercedes.com. Similarly, internet merchants may use trademarks to unite even non-proximate markets from the consumer's perspective. Suppose a seller markets MERCEDES bubble gum. In the brick-and-mortar world, a buyer is relatively unlikely to encounter the product while shopping for a MERCEDES automobile. Online, however, that sort of encounter is far more likely thanks to search engines and other algorithm-driven devices that may direct traffic to related sites and advertisements.

It is also possible that the internet increases consumer exposure to confusing marks by increasing the velocity with which one moves between markets. A consumer walking through a mall will encounter a variety of stores and their affiliated marks. But the total range will be constrained both by the kinds of stores that tend to populate the malls and the speed with which the shopper can move around before needing to go elsewhere. Both constraints are reduced online, potentially increasing consumer contact with similar marks in different markets. One is more likely to bump into the marks for DELTA faucets and DELTA airlines online than when walking through a mall.

Third, as alluded to above, the internet has introduced new mechanisms for similar marks to come into contact. For example, a search for "pandora" has as its top two hits the page for the online radio station and jewelry store. Or, in the case of keyword advertising, use of a trademarked term may provoke the collision between competing marks in display advertising accompanying search results.

Fourth, and relatedly, the internet generated new ways to use trademarks, creating new avenues for the (alleged) weakening of mark meaning, like the use of trademarks in domain names and keywords discussed in greater detail below.

2. *The history*

At the turn of this century, lawmakers expanded trademark's domain in a manner consistent with the logic described above: If the internet exposes more consumers to potentially misleading mark uses, then trademark powers must be strengthened to protect the primacy of the trademark holder's version of the mark. In the late 1990s and early 2000s courts experimented with a range of doctrines to chase the novel trademark uses appearing online. Judges invoked dilution theories;¹¹⁰ they recalibrated the multifactor test for online cases;¹¹¹ and they invigorated the largely dormant doctrine of initial interest confusion.¹¹² In the domain name arena, public and private legislation stepped in to take some pressure off the courts. Congress passed the Anti-cybersquatting Consumer Protection Act ("ACPA"), and the Internet Corporation for Assigned Names and Numbers ("ICANN") instituted the Uniform

outlived its usefulness.").

¹¹⁰ See, e.g., *Panavision Int'l, L.P. v. Toeppen*, 141 F.3d 1316, 1319 (9th Cir. 1998).

¹¹¹ See, e.g., *GoTo.com, Inc. v. Walt Disney Co.*, 202 F.3d 1199, 1205 (9th Cir. 2000).

¹¹² See, e.g., *Brookfield Commc'ns, Inc. v. W. Coast Entm't Corp.*, 174 F.3d 1036, 1062-63 (9th Cir. 1999).

Domain Name Dispute Resolution Policy (“UDRP”), which funneled trademark-based domain name disputes into the private dispute resolution process.¹¹³

Many of these cases rested, in part, on the concern that consumers would encounter trademarks in unfamiliar settings without the distinguishing context consumers rely on in more familiar markets.¹¹⁴ Because the internet collapses market boundaries, this context could not be relied upon as it had been in the past.¹¹⁵

That logic seems to counsel continued growth for trademark rights if stability of meaning is what matters. The internet puts us on a slope whose endpoint is in essence powerful trademarks in one (global) market. But courts have not gone so far. Their reluctance fits another story we can tell about the relationship between the internet and trademark information.

ii. The internet and the retreat of trademark authority

I. The story

We might instead view the internet as a direct attack on the trademark model of authority. On this view, the internet’s ability to undermine expert gatekeepers and enable alternative communities of meaning, each with their own incompatible “truths,”¹¹⁶ extends to trademarks. For example, some trademarks have different owners in different nations. If the internet truly collapses markets in a way that requires trademark exclusivity, this should be untenable. But of course it isn’t. Hasbro owns the SCRABBLE mark in the United States and Canada; a subsidiary of Mattel owns it in the rest of the world.¹¹⁷ The internet lets consumers become aware of these incompatible uses and, with filtering, they may choose an extraterritorial definition of SCRABBLE. As it is, SCRABBLE-related web pages of both companies are accessible online.¹¹⁸ Life goes on. Internet users are capable of

¹¹³ Michael Grynberg, *More Than IP: Trademark Among the Consumer Information Laws*, 55 WM. & MARY L. REV. 1429, 1485–86 (2014).

¹¹⁴ See, e.g., *Playboy Enters., Inc. v. Netscape Commc’ns Corp.*, 354 F.3d 1020, 1025 (9th Cir. 2004) (arguing that online ads keyed to trademark terms might constitute infringement even if consumers immediately realize upon landing on the linked website that it is unrelated to the plaintiff); *GoTo.com, Inc.*, 202 F.3d at 1206 (“Whereas in the world of bricks and mortar, one may be able to distinguish easily between an expensive restaurant in New York and a mediocre one in Los Angeles, the Web is a very different world.” (citation omitted)).

¹¹⁵ See *GoTo.com, Inc.*, 202 F.3d at 1205 (“In the context of the Web in particular, the three most important *Sleekcraft* factors are (1) the similarity of the marks, (2) the relatedness of the goods or services, and (3) the ‘simultaneous use of the Web as a marketing channel.’” (citation omitted)).

¹¹⁶ See *supra* Section III.A.

¹¹⁷ *Welcome to the Official Worldwide SCRABBLE Home Page*, SCRABBLE, www.scrabble.com [<https://perma.cc/B2K2-UBD6>] (“SCRABBLE® is a registered trademark. All intellectual property rights in and to the game are owned in the U.S.A and Canada by Hasbro Inc., and throughout the rest of the world by J.W. Spear & Sons Limited of Maidenhead, Berkshire, England, a subsidiary of Mattel Inc. Mattel and Spear are not affiliated with Hasbro.”).

¹¹⁸ *Compare Scrabble*, MATEL GAMES, <http://www.mattelgames.com/en-my/scrabble> [<https://perma.cc/U8X2-KX5U>] (showing the Mattel Scrabble website for the game available outside of the United States and Canada), *with SCRABBLE*, <https://scrabble.hasbro.com/en-us> [<https://perma.cc/S6Z9-RGJZ>] (showing the Hasbro Scrabble website for the game available in the United States and Canada).

segmenting themselves into belief communities that dispute the curvature of the earth; the ability to vary trademark definitions seems trivial in comparison.¹¹⁹

Moreover, the accumulation and persistence of data online likely make this state of affairs inevitable. No amount of trademark policing will remove all references to alternative mark uses. Even if, for example, the European Union were to convince the United States to grant certification mark rights to “parmesan” cheese,¹²⁰ the existing generic uses of the term would likely linger online long after store shelves come into compliance. This would also be true of infringing uses, for stamping them out not only requires removing all directly infringing uses, which is difficult, but also any references to them, which is impossible if the referrals are not themselves actionable. To the extent nothing is forgotten online, fodder for alternative trademark narratives will likely always be present. While this may undermine trademark reliability in some cases, it also unlocks positive information externalities by allowing sellers to communicate brand similarity where such uses do not create a risk of source confusion at the point of sale.¹²¹

2. *The history*

Early online trademark cases were notorious for treating consumers as incapable of appreciating the context of certain novel trademark uses. For example, defendants faced liability for attempts to be listed in a search engine result for trademarked terms.¹²² The fact that the defendant website was clearly distinguishable from the plaintiffs’ often did not matter.¹²³ The harm, such as it was, was the potential errant click from a search results page.¹²⁴

Though the logic of these cases rested to a large extent on free-riding considerations, they treated the act of arriving on the “wrong” site as something beyond the consumer’s control and, therefore, responsibility.¹²⁵ Over time, however, courts have gradually become more appreciative of consumers’ ability to consider context.¹²⁶ They are increasingly likely to see internet users as more sophisticated as a normative matter, regardless of whether this is true empirically.¹²⁷

The shift can be seen in the Ninth Circuit’s treatment of *Amazon.com*,¹²⁸ which in some ways might be seen as a forerunner of the hypothesized AI of the thought

¹¹⁹ In the *Scrabble* case, of course, the companies are marketing a game that is largely, though not completely, identical. In other cases, a similar mark name might involve greater differences from nation to nation. See generally, e.g., *Harrods Ltd. v. Sixty Internet Domain Names*, 302 F.3d 214 (4th Cir. 2002) (discussing two companies with the same rights to use “Harrods” in separate geographic regions).

¹²⁰ In the United States, many cheeses claim to be “parmesan” but the certification mark for the region-specific product is PARMIGIANO-REGGIANO. See PARMIGIANO-REGGIANO, Registration No. 1,896,683.

¹²¹ Grynberg, *supra* note 27, at 62 n.215, 62–63.

¹²² See, e.g., *Brookfield Commc’ns, Inc. v. W. Coast Entm’t Corp.*, 174 F.3d 1036, 1062–67 (9th Cir. 1999).

¹²³ See, e.g., *id.* at 1062.

¹²⁴ See *id.*

¹²⁵ See generally Michael Grynberg, *The Road Not Taken: Initial Interest Confusion, Consumer Search Costs, and the Challenge of the Internet*, 28 SEATTLE U. L. REV. 97 (2004) (providing an analysis of case law).

¹²⁶ See, e.g., *Toyota Motor Sales, U.S.A., Inc. v. Tabari*, 610 F.3d 1171, 1178–79 (9th Cir. 2010).

¹²⁷ Compare *id.* at 1178 (contending that consumers who shop on the internet are generally sophisticated), with *id.* at 1185–86 (Fernandez, J., concurring) (disputing contention).

¹²⁸ See *Multi Time Mach., Inc. v. Amazon.com, Inc.*, 804 F.3d 930, 933 (9th Cir. 2015) (issuing a

experiment. It is highly algorithmic; it functions in large part by “knowing” the consumer; and it limits the importance of traditional trademarks and brands. The next Part considers this claim in greater detail.

IV. NASCENT AIS

Various forms of artificial intelligence already influence consumer choices online.¹²⁹ We may be a long way off from the AI of the hypothetical, but trademark law is already addressing a world of machine-guided purchasing.

In this, Amazon.com provides a helpful illustration of many of the trademark issues described above. As Amazon customers know, the website guides consumer purchases by providing search suggestions that draw on a massive database of product information as well as individual consumer data, reflected by past purchases and browsing history. Amazon is also working hard to make its AI technology a more active participant in consumer purchasing decisions.¹³⁰

judgment in favor of Amazon.com upon finding that “no reasonably prudent consumer accustomed to shopping online would [] be confused as to the source of the products” in Amazon’s search results).

¹²⁹ See, e.g., Blake Morgan, *How Amazon has Reorganized Around Artificial Intelligence and Machine Learning*, FORBES (July 16, 2018, 2:37 PM), <https://www.forbes.com/sites/blakemorgan/2018/07/16/how-amazon-has-re-organized-around-artificial-intelligence-and-machine-learning/#252ad94d7361> [<https://perma.cc/M7WW-PPTE>] (“AI also plays a huge role in Amazon’s recommendation engine, which generates 35% of the company’s revenue. Using data from individual customer preferences and purchases, browsing history and items that are related and regularly bought together, Amazon can create a personalized list of products that customers actually want to buy.”). And there is no shortage of breathless predictions surrounding the purported benefits. See Paul Roetzer, *This AI Tool Gets Retail Customers to Buy More—and Get Smarter Over Time*, MARKETING ARTIFICIAL INTELLIGENCE INST. (July 12, 2018), <https://www.marketinginstitute.com/blog/this-ai-tool-gets-retail-customers-to-buy-more-and-get-smarter-over-time> [<https://perma.cc/E8BB-EQ5J>] (advertising a particular technology); see also *The Future of Artificial Intelligence in Consumer Experience: According to the AT&T Foundry*, ROCKETSPACE, <https://www.rocketSPACE.com/hubfs/accelerator/the-future-of-artificial-intelligence.pdf?hsLang=en-us> [<https://perma.cc/3S5C-MKXG>] (“Due to a deep comprehension of the customer, brands will provide sublime experiences catered to users’ behavioral patterns. Everything from shopping to driving will draw from user behavior to become highly pertinent and personalized to the end consumer. Intelligent prediction and optimization will allow the consumer to feel that each branded product or experience is made just for them.”).

¹³⁰ As Jeff Bezos stated in his 2017 letter to shareholders:

At Amazon, we’ve been engaged in the practical application of machine learning for many years now. Some of this work is highly visible: . . . [including] Alexa, our cloud-based AI assistant. . . .

But much of what we do with machine learning happens beneath the surface. Machine learning drives our algorithms for demand forecasting, product search ranking, product and deals recommendations, merchandising placements, fraud detection, translations, and much more. Though less visible, much of the impact of machine learning will be of this type – quietly but meaningfully improving core operations.

Taylor Soper, *Full Text: In Annual Shareholder Letter, Jeff Bezos Explains Why it will Never be Day 2 at Amazon*, GEEKWIRE (Apr. 12, 2017, 8:22 AM), <https://www.geekwire.com/2017/full-text-annual-letter-amazon-ceo-jeff-bezos-explains-avoid-becoming-day-2-company/> [<https://perma.cc/UVV9-UKM6>].

Consumers naturally use the site to search for specific trademarked goods.¹³¹ But as discussed in the following subparts, Amazon.com also shows three ways that the coming world may deemphasize trademarks and trademark law. First, the site uses trademark information to generate alternative options to searched-for products.¹³² This practice has survived a legal challenge, reflecting the increasing judicial comfort with limiting trademark holder authority online.¹³³ Second, the Amazon website limits the importance of trademarks.¹³⁴ For many, the Amazon platform matters more than the marks found on it. Those consumers who prefer the Amazon ecosystem to shopping elsewhere appear to pay less attention to trademark information than the traditional story of trademark law would suggest. Their actions are more in line with the world of the hypothesized AI. Third, the platform presents more than its share of problems as a model AI (particularly its lack of loyalty to its customers), but these problems are not trademark problems.¹³⁵

A. Amazon and trademark's online retreat

As discussed in Part II, developments in AI may reduce the importance of trademarks to ultimate purchasing decisions.¹³⁶ Doctrinally, this would mean that consumers (or their avatars) will be expected to assume a greater role in assessing the context of third-party trademark uses, rather than letting trademark holders control them.

*Multi Time Machine, Inc. v. Amazon.com, Inc.*¹³⁷ takes a large step in this direction. The watchmaker Multi Time Machine (“MTM”) objected to Amazon.com’s use of its marks to arrange information for site users.¹³⁸ Amazon does not carry MTM watches.¹³⁹ When Amazon received search queries for the “MTM special ops” brand, it still returned a results list that included competing brands.¹⁴⁰ Amazon did not simply advise the searcher of the brand’s absence from Amazon.¹⁴¹

All of the search results were labeled accurately, but MTM nonetheless alleged potential initial interest confusion.¹⁴² The purported danger was that consumers might speculate about a connection between the searched-for trademark and the

¹³¹ Amazon rivals Google with respect to use for product search. Krista Garcia, *More Product Searches Start on Amazon*, EMARKETER (Sept. 7, 2018), <https://www.emarketer.com/content/more-product-searches-start-on-amazon> [https://perma.cc/C9DQ-ZPXS] (“A number of consumer surveys have shown that more US digital shoppers now start their searches on Amazon. Nearly half (46.7%) of US internet users started product searches on Amazon compared with 34.6% who went to Google first, according to a May 2018 Adeptmind survey.”).

¹³² See *infra* Section IV.A.

¹³³ See *infra* Section IV.A.

¹³⁴ See *infra* Section IV.B.

¹³⁵ See *infra* Section IV.C.

¹³⁶ See *supra* Part II.

¹³⁷ 804 F.3d 930 (9th Cir. 2015).

¹³⁸ *Id.* at 932–33.

¹³⁹ *Id.* at 932.

¹⁴⁰ *Id.* at 932–33.

¹⁴¹ *Id.* at 936.

¹⁴² *Id.*

returned results.¹⁴³ The dissent treated the concern as plausible.¹⁴⁴ Nonetheless, the majority imputed a higher level of sophistication on the part of Amazon.com users:

MTM argues that initial interest confusion might occur because Amazon lists the search term used—here the trademarked phrase “mtm special ops”—three times at the top of the search page. MTM argues that because Amazon lists the search term “mtm special ops” at the top of the page, a consumer might conclude that the products displayed are types of MTM watches. But, merely looking at Amazon’s search results page shows that such consumer confusion is highly unlikely. None of these watches is labeled with the word “MTM” or the phrase “Special Ops,” let alone the specific phrase “MTM Special Ops.” . . . [N]o reasonably prudent consumer accustomed to shopping online would view Amazon’s search results page and conclude that the products offered are MTM watches. It is possible that someone, somewhere might be confused by the search results page. But, “[u]nreasonable, imprudent and inexperienced web-shoppers are not relevant.”¹⁴⁵

In the court’s view, internet users are capable of reading information in context. The majority trusted (and expected) consumers to understand that the returned results meant “these are results *similar* to the watch brand you entered” and not “these are results *sponsored by* (or affiliated with) the watch brand you entered.” In making this assumption, the court allowed Amazon and its users to take advantage of the information externalities of the MTM mark. These spillovers provided Amazon with an efficient mechanism for communicating the existence of alternatives to MTM’s product.

But what of the “costs” of the activity? Even if there is no risk of source confusion, given the accurate labeling of the search results, some consumers are looking for a particular brand and *only* that brand. Isn’t there a cost to making them cut through the clutter to find the branded result? Here, the opinion accepts that trademark law can only do so much. The nature of the internet means that there always will be clutter that accompanies any search. Consumers would still face the task of sorting results even if MTM had prevailed:

Further, some of the products listed are not even watches. The sixth result

¹⁴³ *Id.* at 933, 938. This purported danger finds some support in other cases. *See, e.g.*, *Maker’s Mark Distillery, Inc. v. Diageo North America, Inc.*, 679 F.3d 410, 424 (6th Cir. 2012); *Brookfield Commc’ns, Inc. v. W. Coast Entm’t Corp.*, 174 F.3d 1036, 1057 (9th Cir. 1999); *see also supra* text accompanying notes 39–46.

¹⁴⁴ *Multi Time Mach., Inc.*, 804 F.3d at 940 (Bea, J., dissenting) (“Because I believe that an Amazon shopper seeking an MTM watch might well initially think that the watches Amazon offers for sale when he searches ‘MTM Special Ops’ are affiliated with MTM, I must dissent.”). Indeed, Judge Bea initially authored a majority opinion in MTM’s favor, but it was superseded and replaced by the opinion discussed in the text. *Multi Time Mach., Inc. v. Amazon.com, Inc.*, 792 F.3d 1070 (9th Cir. 2015), *withdrawn and superseded on reh’g*, 804 F.3d 930 (9th Cir. 2015).

¹⁴⁵ *Multi Time Mach., Inc.*, 804 F.3d at 938 (majority opinion) (quoting *Toyota Motor Sales, U.S.A., Inc. v. Tabari*, 610 F.3d 1171, 1176 (9th Cir. 2010)).

is a book entitled “**Survive!: The Disaster, Crisis and Emergency Handbook** by Jerry Ahem.” The tenth result is a book entitled “**The Moses Expedition: A Novel** by Juan Gómez-Jurado.” No reasonably prudent consumer, accustomed to shopping online or not, would assume that a book entitled “The Moses Expedition” is a type of MTM watch or is in any way affiliated with MTM watches.¹⁴⁶

This is the filtration information model at work. All sorts of data are out there; this clutter requires the use of search tools to sift and categorize information. Our choices may or may not be compatible with the desires of trademark holders, but the markholders cannot dictate them. In essence, the majority accepts that “mtm special ops” is a term around which multiple communities of meaning may arise. Consumers have both the freedom to sort the data surrounding the term and the obligation to own the task. The trademark holder has no privileged authority beyond a very narrow class of meanings (in this case, watches specifically branded MTM SPECIAL OPS).

Multi Time Machine’s embrace of context echoes other cases. In *Network Automation, Inc. v. Advanced Systems Concepts, Inc.*, the Ninth Circuit considered the issue of trademark-triggered advertising, concluding that such advertising does not create liability unless the displayed advertisement *itself* creates likely confusion.¹⁴⁷ Specifically, the court modified the use of the multifactor test in keyword advertising cases:

[T]he most relevant factors to the analysis of the likelihood of confusion are: (1) the strength of the mark; (2) the evidence of actual confusion; (3) the type of goods and degree of care likely to be exercised by the purchaser; and (4) the labeling and appearance of the advertisements and the surrounding context on the screen displaying the results page.¹⁴⁸

This fourth factor, absent from the traditional multifactor test,¹⁴⁹ effectively brings false advertising thinking into the case.

The result echoes a point made above: False advertising doctrine is often a better option than trademark law for addressing certain allegedly misleading activities involving trademarks online.¹⁵⁰ And, as did *Multi Time Machine*, the opinion accepts that the online world requires consumers to sort through context without automatically making assumptions based on the mere presence of a trademark. If there is liability to be had by pairing information with a mark, there should be something actionable about that information as reflected by advertising and other “surrounding context.”¹⁵¹

Domain names provide a final example of the liberalization of online use of trademarks. *Toyota Motor Sales, U.S.A., Inc. v. Tabari* uses the nominative fair use

¹⁴⁶ *Id.*

¹⁴⁷ See *Network Automation, Inc. v. Advanced Sys. Concepts, Inc.*, 638 F.3d 1137, 1149, 1154 (9th Cir. 2011).

¹⁴⁸ *Id.* at 1154.

¹⁴⁹ See *id.* at 1153–54.

¹⁵⁰ See *supra* Section II.B.3.

¹⁵¹ See *Network Automation*, 638 F.3d at 1154.

doctrine to allow the unauthorized use of trademarks in third-party domain names.¹⁵² The opinion sees consumers as knowing better than to attach too much importance to the appearance of a trademark in a domain name when the URL contains other information:

[Consumers] fully expect to find some sites that aren't what they imagine based on a glance at the domain name or search engine summary. Outside the special case of trademark.com, or domains that actively claim affiliation with the trademark holder, consumers don't form any firm expectations about the sponsorship of a website until they've seen the landing page—if then. This is sensible agnosticism, not consumer confusion.¹⁵³

The opinion also appreciates the range of purposes that might surround the use of a mark within a domain name:

But the case where the URL consists of nothing but a trademark followed by a suffix like .com or .org is a special one indeed. The importance ascribed to trademark.com in fact suggests that far less confusion will result when a domain making nominative use of a trademark includes characters in addition to those making up the mark. Because the official Lexus site is almost certain to be found at lexus.com (as, in fact, it is), it's far less likely to be found at other sites containing the word Lexus. On the other hand, a number of sites make nominative use of trademarks in their domains but are not sponsored or endorsed by the trademark holder: You can preen about your Mercedes at mercedesforum.com and mercedestalk.net, read the latest about your double-skim-no-whip latte at starbucks gossip.com and find out what goodies the world's greatest electronics store has on sale this week at fryselectronics-ads.com. Consumers who use the internet for shopping are generally quite sophisticated about such matters and won't be fooled into thinking that the prestigious German car manufacturer sells boots at mercedesboots.com, or homes at mercedeshomes.com, or that comcastsucks.org is sponsored or endorsed by the TV cable company just because the string of letters making up its trademark appears in the domain.¹⁵⁴

In other words, different online communities of meaning may grow up around trademarks without harm to consumers.

B. One brand to rule them all?

As discussed above, the information-simplification offered by trademarks is a double-edged sword. When consumers shop for favored brands that satisfy

¹⁵² Toyota Motor Sales, U.S.A., Inc. v. Tabari, 610 F.3d 1171, 1175–77 (9th Cir. 2010).

¹⁵³ *Id.* at 1179 (citation omitted).

¹⁵⁴ *Id.* at 1178 (citations omitted).

established preferences, they avoid a costlier search process that might lead them to better results.¹⁵⁵ The hypothesized AI offers a way out by doing the work of compiling the data that would identify better products.¹⁵⁶ The resulting marketplace would deemphasize trademarks in important respects.¹⁵⁷

But today's second-best world shows another way AIs may deemphasize trademarks. As illustrated by *Multi Time Machine*, many Amazon users engage in another form of information simplification. By shopping at Amazon.com, they limit their search to the Amazon ecosystem. That choice sets the range of available trademarks. Rather than searching by selecting from among trademarks, many Amazon users select Amazon, relying on its algorithms to narrow the range of possibilities. To the extent particular brand names are not Amazon offerings, they are excluded.¹⁵⁸ They remain available elsewhere, but they are pre-filtered from consideration. Stated another way, there is trademark work going on, but it applies to platform competition (the selection of Amazon) and not necessarily purchasing decisions beyond that (what is bought on Amazon).¹⁵⁹

¹⁵⁵ See *supra* Section II.A.

¹⁵⁶ See *supra* Sections I.B, II.A.

¹⁵⁷ See *supra* Section II.A.

¹⁵⁸ Indeed, this selection effect is particularly strong for Amazon Prime users:

Amazon Prime users are both more likely to buy on its platform and less likely to shop elsewhere. “[Sixty-three percent] of Amazon Prime members carry out a paid transaction on the site in the same visit,” compared to 13% of non-Prime members. For Walmart and Target, those figures are 5% and 2% respectively. One study found that less than 1% of Amazon Prime members are likely to consider competitor retail sites in the same shopping session.

Lina M. Khan, *Amazon's Antitrust Paradox*, 126 YALE L.J. 710, 752 (2017) (footnotes omitted) (quoting Clare O'Connor, *Walmart and Target Being Crowded Out Online by Amazon Prime*, FORBES (Apr. 6, 2015, 12:59 PM), <http://www.forbes.com/sites/clareoconnor/2015/04/06/walmart-and-target-being-crowded-out-online-by-amazon-prime> [<https://perma.cc/X6WC-6C7J>]).

¹⁵⁹ To be sure, this kind of information masking exists in traditional offline purchase decisions, as when consumers choose to shop at a Wal-Mart or Target, and those choices allow consumers to offload some amount of search costs to the retailer:

Retailers also reduce manufacturer-consumer transaction costs by satisfying the needs of consumer niche markets more cheaply than manufacturers can. Determining consumer needs is costly; it includes costs to aggregate data, analyze it, and respond to identified needs. For manufacturers trying to cater to multiple, diverse, and far-flung consumer segments, it can be cost-prohibitive to learn and understand the needs of every consumer niche, especially small niches. In contrast, retailers can cater to consumer niches, such as specific geographies or industries.

Goldman, *supra* note 50, at 413. But online retailers may have features—the ability to mine data at the individual level, a greater breadth of possible offerings, and an ability to offer tailored nudges—that may make online retailers different in kind:

Retailers often leverage their power and custody of the consumer to swap out brands for their own private label. That's nothing new. Only we've never seen any retailer this good at it. Amazon, armed with infinite capital provided by eager investors, is leading a war on brands to starch the margin from brands and deliver it back to the consumer.

SCOTT GALLOWAY, *THE FOUR: THE HIDDEN DNA OF AMAZON, APPLE, FACEBOOK, AND GOOGLE* 51–52 (2017); cf. Khan, *supra* note 158, at 782 (noting that while traditional retailer stores sometimes use house brands to compete in vertical markets, such stores “are generally only able to collect information on actual sales, [while] Amazon tracks

Amazon’s power here is such that it has supplanted Google as the leading starting point for product search,¹⁶⁰ notwithstanding Google’s comparative neutrality with respect to product selection (as it is not involved with retail sales and fulfillment to the extent that Amazon is).¹⁶¹

The phenomenon of Amazon’s increased use of its own and affiliated brands suggests the vulnerability of trademarks to platform selection. Many Amazon.com searches direct users to affiliated brands with little renown or goodwill outside of Amazon.¹⁶² The success of the practice may illustrate an antitrust issue,¹⁶³ but it also highlights the ability of a consumer agent to negate the power of brands. In the case of Amazon, it allows affiliated brands to develop what looks like near-instantaneous goodwill. Indeed, Amazon’s power is so strong that many of these upstart marks break basic rules of brand attractiveness but nonetheless find their place in the Amazon universe. For example:

A search for “three piece suit” on Amazon returns a litany of budget brands like YFFUSHI, WULFUL and WEEN CHARM. Ungainly names aside, some of these labels have been positively reviewed, overcoming the considerable challenges of branding, and marketing, to an audience thousands of miles away, and sometimes relying on the Amazon seller marketplace and using the company to handle fulfillment — warehousing and shipping, basically.¹⁶⁴

For this to happen, trademarks cannot be performing as much work as the traditional trademark story would suggest. Traditional trademark uses (which brand of pants should I buy on Amazon.com?) are simply less important in the Amazon-mediated environment (which pants does Amazon suggest?).

It is entirely possible, of course, that in the near term, the net effect of online technology is to raise, rather than lower, the importance of strong trademarks. If the internet makes available more marks in a product category than could ever be found in a store, then looking to well-known marks is one way to manage the prospect of information overload.¹⁶⁵ Sure.¹⁶⁶ But there is no reason to assume that consumers

what shoppers are searching for but cannot find, as well as which products they repeatedly return to, what they keep in their shopping basket, and what their mouse hovers over on the screen”).

¹⁶⁰ Garcia, *supra* note 131.

¹⁶¹ Cf. John Herman, *Everything on Amazon is Amazon!*, N.Y. TIMES (Nov. 15, 2018), <https://www.nytimes.com/2018/11/15/style/this-is-also-amazon.html> [<https://perma.cc/T5DL-PDAD>].

¹⁶² See *id.* (“There are vanishingly few types of consumer goods that you can’t buy, in some form, on Amazon. But it is missing plenty of *brands*. In 2009, the company started selling products under its own name. It soon moved beyond the first AmazonBasics—items including budget electronics and batteries—to a wider range of Amazon-branded products. This was followed by an explosion of company-owned brands, including dozens with Amazon-free names.”).

¹⁶³ Khan, *supra* note 158, at 780–83.

¹⁶⁴ Herman, *supra* note 161.

¹⁶⁵ See Jacob Jacoby et al., *Brand Choice Behavior as a Function of Information Load*, 11 J. MARKETING RES. 63, 68 (1974); see also Goldman, *supra* note 50, at 414–16.

¹⁶⁶ But see *supra* note 159 and accompanying text (discussing online retailers’ ability to manage information for consumers); *infra* note 168 and accompanying text.

will forever be making the final decision themselves. Designating an agent to compile a list of options is just a step along the way to trusting it to choose from among the list or, as already happens, being open to its nudges.¹⁶⁷ And in any case, there is already evidence that brands occupy less of our mental space than they did in the past.¹⁶⁸

C. Controlling avatars

This Article uses a conception of AI to examine trademark law, not vice versa. That said, thinking of Amazon.com as an AI forerunner also illustrates some potential dangers of offloading decisions to our digital assistants.

First, of course, the Amazon AI is a far cry from the thought experiment in terms of capabilities. These limits may have trademark effects in simple ways. For example, Amazon programming often recommends products based on past purchases,¹⁶⁹ thus potentially exacerbating the problem of consumer inertia (where future, better technology might ameliorate it).¹⁷⁰

Second, Amazon's business practices raise concerns regarding the competitive marketplace. These exist with respect to its marketing and shipping infrastructure, but also the data Amazon has about its customers, which gives Amazon a competitive advantage.¹⁷¹ A similar issue exists with respect to whether consumer-focused AI needs access to data about general consumer preferences in order to be effective in assisting particular consumers. Competitive considerations aside, privacy concerns may complicate compiling such data.¹⁷²

Third, more powerful AIs may be used to manipulate consumers into suboptimal behavior,¹⁷³ be it purchasing second-best goods, engaging in needless transactions, or simply continuing to generate valuable data for the creation of other products.¹⁷⁴ Likewise, the AI may systematically direct consumers to purchases selected to serve someone else's needs. Rather than a bespoke world of goods and services tailored precisely to individual consumer preferences, the marketplace may simply be one of a few uber-brands using AI to serve their needs, manufacturing consumer preferences as necessary. Aspects of that problem already confront us, as the technologies that were touted as liberating or empowering individuals now manipulate them on a scale (in terms of the number of people receiving individually

¹⁶⁷ These nudges may be easier to implement when we order by voice rather than interacting with a screen. GALLOWAY, *supra* note 159, at 51 ("In key categories like batteries, Alexa will suggest Amazon Basics, their private label, and play dumb about other choices ('Sorry, that's all I found!') when there are several other brands on amazon.com.").

¹⁶⁸ *Id.* at 48 ("In 2004, 47 percent of affluent consumers could name a favorite retail brand; six years later that number dropped to 28 percent."); *id.* at 50 (noting a declining "percentage of affluents who can identify a 'favorite brand'" in the categories of fashion, jewelry, luxury hotels, and retailer).

¹⁶⁹ Morgan, *supra* note 129.

¹⁷⁰ See *supra* text accompanying notes 33–35.

¹⁷¹ See Khan, *supra* note 158, at 782–83.

¹⁷² See generally ZUBOFF, *supra* note 18.

¹⁷³ See, e.g., Calo, *supra* note 19, at 1021 ("[D]igital market manipulation combines, for the first time, a certain kind of *personalization* with the intense *systemization* made possible by mediated consumption.").

¹⁷⁴ ZUBOFF, *supra* note 18, at 377, 451.

calibrated nudges) never conceived of before the advent of the internet.¹⁷⁵ Engineering machines to do so is itself an academic discipline.¹⁷⁶

This problem is not, however, one of trademark, but rather loyalty.¹⁷⁷ The Amazon AI and those to follow in its footsteps are suspect because they serve someone other than the consumers they purport to assist.¹⁷⁸ The dilemma may not be

¹⁷⁵ See, e.g., ROGER MCNAMEE, ZUCKED: WAKING UP TO THE FACEBOOK CATASTROPHE 85 (2019) (“The artificial intelligences of companies like Facebook (and Google) now include behavioral prediction engines that anticipate our thoughts and emotions, based on patterns found in the reservoir of data they have accumulated about users. Years of Likes, posts, shares, comments, and Groups have taught Facebook’s AI how to monopolize our attention.”); O’NEIL, *supra* note 17, at 183–85 (describing experiments in how Facebook, Google, and other search engines may manipulate users); Jon Brooks, *Tech Insiders Call Out Facebook for Literally Manipulating Your Brain*, KQED (May 25, 2017), <https://www.kqed.org/futureofyou/379828/tech-insiders-call-out-facebook-for-literally-manipulating-your-brain> [<https://perma.cc/LBJ6-DHWX>]; Vinu Goel, *Facebook Tinkers with Users’ Emotions in News Feed Experiment, Stirring Outcry*, N.Y. TIMES (June 29, 2014), <https://www.nytimes.com/2014/06/30/technology/facebook-tinkers-with-users-emotions-in-news-feed-experiment-stirring-outcry.html> [<https://perma.cc/7T8X-CRE8>]; Trevor Haynes, *Dopamine, Smartphones & You: A Battle for Your Time*, HARVARD U.: SCI. NEWS (May 1, 2018), <http://sitn.hms.harvard.edu/flash/2018/dopamine-smartphones-battle-time/> [<https://perma.cc/NL23-5BN3>]; Hope Reese, *Break Up with Your Smartphone*, VOX (Feb. 9, 2018, 10:10 AM), <https://www.vox.com/conversations/2018/2/9/16994794/smartphone-tech-addiction> [<https://perma.cc/MJD4-NFVV>]; James Vincent, *Former Facebook Exec Says Social Media is Ripping Apart Society*, VERGE (Dec. 11, 2017, 6:07 AM), <https://www.theverge.com/2017/12/11/16761016/former-facebook-exec-ripping-apart-society> [<https://perma.cc/BJ2X-RB7H>].

¹⁷⁶ B.J. FOGG, PERSUASIVE TECHNOLOGY: USING COMPUTERS TO CHANGE WHAT WE THINK AND DO 5 (2003) (“[C]aptology focuses on the design, research, and analysis of interactive computing products created for the purpose of changing people’s attitudes or behaviors. It describes the area where technology and persuasion overlap.” (citation omitted)). For an account of how these techniques may be used, see NIR EYAL WITH RYAN HOOVER, HOOKED: HOW TO BUILD HABIT-FORMING PRODUCTS 179–80 (2014).

¹⁷⁷ Jack M. Balkin, *Information Fiduciaries and the First Amendment*, 49 U.C. DAVIS L. REV. 1183, 1227 (2016) (“[O]nline service providers present the familiar problems that generally give rise to fiduciary obligations. First, there are significant asymmetries of knowledge and information between online service providers and end-users. Second, it is very difficult for end-users to verify online companies’ representations about data collection, security, use, and dissemination. Third, it is very difficult for end-users to understand what online companies do with their data and how data analysis and use affects their interests. Fourth, even if end-users understood these information practices, it would be almost impossible for end-users to monitor them.”).

¹⁷⁸ As observed about the—now quaint—problems of targeted advertising, current technology is at odds with the promise of AIs that would serve consumer needs:

The idealists foresaw a day when ad platforms would be like a loyal valet who detected his master’s needs before he was aware of them, who suggested a new pair of shoes as a reasonably priced replacement for those you hadn’t noticed were wearing out. Perhaps he would remind you of your mother-in-law’s birthday while offering to send an appropriate gift at a one-day discount.

But the gap between this theory and its execution was wide enough to march Kitchener’s Army through it. Google’s CEO Eric Schmidt had once said that the ideal was to “get right up to the creepy line and not cross it.” Unfortunately, by the mid-2010s, that line was being crossed constantly. While promising to be “helpful” or “thoughtful,” what was delivered was often experienced as “intrusive” and worse. Some ads seemed more like stalkers than valets: if, say, you’d been looking at a pair of shoes on Amazon, an ad for just those shoes would begin following you around the web, prodding you to take another look at them.

TIM WU, THE ATTENTION MERCHANTS: THE EPIC SCRAMBLE TO GET INSIDE OUR HEADS 323 (2016) (footnotes omitted). Whether the tailored ads are hidden manipulation or seem more like overt harassment may to some extent be a question of technical sophistication. As Wu observes, many ads:

satisfied by competition, particularly if top-of-the-line AI proves to be the product of a winner-take-all market or if the data necessary to make it effective is a byproduct of a dominant share of some other market. Perhaps the issue can be addressed by antitrust law, but it may be necessary to mandate a duty of loyalty for consumer digital avatars. This issue would be, to say the least, complex as it reintroduces the difficult question, ducked above, of what loyalty means given the difficulty of determining what consumer preferences actually are.¹⁷⁹

V. THE LARGER LIMITS OF TRADEMARK

Trademark law's inability to meet these challenges points the way to other doctrines. Perhaps false advertising, antitrust, or consumer protection laws will fill the necessary gaps. Perhaps entirely new doctrines and bodies of law will emerge to fetter AI and promote consumer interests.

Whatever fills the gap, before long the law will have to squarely face the question not only of what consumers want descriptively, but also—because the answer to this question is a byproduct of consumer interactions with AI technology—what they *should* want as a normative matter.

We already face the question in the much-bemoaned problem of online fake news. Fake news illustrates a number of this Article's themes. First, it illustrates the death of authority online.¹⁸⁰ There are no curators who might make these falsehoods hard to find. Instead, they are cheaply created and waiting online for anyone to find and use as they see fit.¹⁸¹

Second, fake news shows the limits of trademarks in the online world (and the increasing importance of platforms¹⁸²). We might have thought that trademark law would be the natural solution to fake news. Given the high volume of falsehoods online, reputable news sites should rise to the top of the information market, limiting the reach of false stories and conspiracy theories.¹⁸³

... turned out to be more of a studied exploitation of one's weaknesses. The overweight were presented with diet aids; the gadget-obsessed plied with the latest doodads; gamblers encouraged to bet; and so on. One man, after receiving a diagnosis of pancreatic cancer, found himself followed everywhere with "insensitive and tasteless" ads for funeral services. The theoretical idea that customers might welcome or enjoy such solicitations increasingly seemed like a bad joke.

Id. at 323–24.

¹⁷⁹ See *supra* note 21.

¹⁸⁰ See *supra* Part III.

¹⁸¹ See, e.g., SANGER, *supra* note 84, at 185 (describing activities of Russian fake news operations during the 2016 Presidential campaign, observing that for "a fraction of the cost for an evening of television advertising on a local American television station. . . . Putin's trolls reached up to 126 million Facebook users, while on Twitter they made 288 million impressions").

¹⁸² See *supra* Part IV.

¹⁸³ Some sites do not mimic any particular page, but rather strive simply to look like a news outlet, and that is enough in many cases. See Abby Ohlheiser, *This is how Facebook's Fake-news Writers Make Money*, WASH. POST (Nov. 18, 2016), https://www.washingtonpost.com/news/the-intersect/wp/2016/11/18/this-is-how-the-internets-fake-news-writers-make-money/?utm_term=.bda32a90ed2a [https://perma.cc/WG82-NFEW]. It is likewise easy to create the appearance of advocacy group activities. For example, during the 2016 campaign, the Internet Research Agency, a Russian organization operating in St. Petersburg, organized a "Stop Islamization of

But trademarks have simply proven to be unequal to the task. Part of the problem is the role of platforms. Sites like Facebook deliver algorithmically selected content, allowing many users to suture themselves into the filter bubble of their choosing,¹⁸⁴ assuming they can be said to be choosing to do so.¹⁸⁵ These sites thus far have little market or regulatory pressure to address the negative externalities of fake news consumption.¹⁸⁶ Moreover, channeling users into simple, extreme categories defined by Left and Right—rather than serving information that would reinforce more moderate tendencies—seems the path of least resistance with current algorithmic technology.¹⁸⁷

Worse, in many cases the problem is not that trademarks are not functioning, but rather that they *are*. In a world in which the president derides unfriendly press as fake, the avoidance of reputable news sites becomes just another form of political expression.¹⁸⁸ Many fake news consumers are getting exactly what they want. The law isn't the problem so much as the citizens it would protect.

That answer carries a tragic dimension for which there may not be law-based answers. In 2018 the *Washington Post* profiled a fake news creator who posts the

Texas” event sponsored by the made-up “Heart of Texas” group. SANGER, *supra* note 84, at 201–02. “Then, in a masterful stroke, the Russians created an opposing group, ‘United Muslims of America,’ which scheduled a counter-rally, under the banner of ‘Save Islamic Knowledge.’ The idea was to motivate actual Americans—who had joined each of the Facebook groups—to face off against each other and prompt a lot of name-calling and, perhaps, some violence.” *Id.* at 202.

¹⁸⁴ There has long been a debate about whether the net effect of the internet is to inform or fence people off into the echo chamber of their choosing. See, e.g., WEINBERGER, *supra* note 2, at 81–83. Both may be possible. Hunt Allcott et al., *The Welfare Effects of Social Media* 1–2 (Nat’l Bureau of Econ. Research, Working Paper No. 25514, 2019), <http://web.stanford.edu/~gentzkow/research/facebook.pdf> [<https://perma.cc/BM6R-35TC>] (study reporting that Facebook deactivation “reduced both factual news knowledge and political polarization” but “increased subjective well-being”).

¹⁸⁵ See MCNAMEE, *supra* note 175, at 92–93 (describing how platforms facilitate extreme views); cf. O’NEIL, *supra* note 17, at 194 (“Successful microtargeting, in part, explains why in 2015 more than 43 percent of Republicans, according to a survey, still believed the lie that President Obama is a Muslim.”).

¹⁸⁶ See, e.g., SANGER, *supra* note 84, at 253–55; see also GALLOWAY, *supra* note 159, at 118 (“[I]f [Facebook] figures out you lean Republican, it will feed you more Republican stuff, until you’re ready for the heavy hitters, the GOP outrage: Breitbart, talk radio clips. You may even get to Alex Jones.”).

¹⁸⁷ See GALLOWAY, *supra* note 159, at 117–19; *id.* at 118 (“Marketing to moderates is like fracking for gas. You only do it if the easier alternatives aren’t available.”).

¹⁸⁸ This is not to say that trademarks are never counterfeited to spread fake news. See, e.g., Ian Stewart, *Real Fake News: Activists Circulate Counterfeit Editions of ‘The Washington Post,’* NPR (Jan. 16, 2019, 1:23 PM), <https://www.npr.org/2019/01/16/685857177/real-fake-news-activists-circulate-counterfeit-editions-of-the-washington-post> [<https://perma.cc/PT56-M2KL>]. But while trademark law can address the problem of sites designed to pass themselves off as a *specific* reputable source, see Joshua Humphrey, *The Plague of Fake News and the Intersection with Trademark Law*, 8 CYBARIS INTELL. PROP. L. REV. 126, 146 (2017) (contending that “trademark infringement can be a successful strategy in combating fake news if likelihood of confusion can be proved”), it has a harder time with those that merely try to look as though they are reputable. Trademark doctrine therefore struggles with activities that create noise that might mask the signal sent by marks that certify quality. See Margaret Chon, *Marks of Rectitude*, 77 FORDHAM L. REV. 2311, 2332, 2343–44 (2009); Grynberg, *supra* note 113, at 1457; cf. 15 U.S.C. § 1127 (2018) (providing that a mark can serve as a trademark “even if th[e] source is unknown”).

most outlandish stories possible as a form of political commentary.¹⁸⁹ His goal is to call attention to the propensity of some on the right wing to believe anything, sure enough:

In the last two years on his page, America's Last Line of Defense, Blair had made up stories about California instituting sharia, former president Bill Clinton becoming a serial killer, undocumented immigrants defacing Mount Rushmore, and former president Barack Obama dodging the Vietnam draft when he was 9. "Share if you're outraged!" his posts often read, and thousands of people on Facebook had clicked "like" and then "share," most of whom did not recognize his posts as satire. Instead, Blair's page had become one of the most popular on Facebook among Trump-supporting conservatives over 55.

"Nothing on this page is real," read one of the 14 disclaimers on Blair's site, and yet in the America of 2018 his stories had become real, reinforcing people's biases, spreading onto Macedonian and Russian fake news sites, amassing an audience of as many 6 million visitors each month who thought his posts were factual. What Blair had first conceived of as an elaborate joke was beginning to reveal something darker. "No matter how racist, how bigoted, how offensive, how obviously fake we get, people keep coming back," Blair once wrote, on his own personal Facebook page. "Where is the edge? Is there ever a point where people realize they're being fed garbage and decide to return to reality?"¹⁹⁰

The story goes on, however, to profile one such believer, and a sadder picture emerges, one less of political activism and more of loneliness and social isolation:

The house was empty and quiet except for the clicking of her computer mouse. She lived alone, and on many days her only personal interaction occurred here, on Facebook. Mixed into her morning news feed were photos and updates from some of her 300 friends, but most items came directly from political groups [she] had chosen to follow: "Free Speech Patriots," "Taking Back America," "Ban Islam," "Trump 2020" and "Rebel Life." Each political page published several posts each day directly into [her] feed, many of which claimed to be "BREAKING NEWS."

...

On display above [her] screen were needlepoints that had once occupied much of her free time, intricate pieces of artwork that took hundreds of hours to complete, but now she didn't have the patience. Out her window was a dead-end road of identical beige-and-brown rock gardens surrounding double-wide trailers that looked similar to her own, many of them occupied by neighbors whom she'd never met. Beyond that was nothing but cactuses and heat waves for as far as she could see — a stretch

¹⁸⁹ Eli Saslow, *'Nothing on this Page is Real': How Lies Become Truth in Online America*, WASH. POST (Nov. 17, 2018, 7:40 PM), https://www.washingtonpost.com/national/nothing-on-this-page-is-real-how-lies-become-truth-in-online-america/2018/11/17/edd44cc8-e85a-11e8-bbdb-72fdbf9d4fed_story.html?noredirect=on&utm_term=.97e60caefcc1 [https://perma.cc/RZG8-EAJB].

¹⁹⁰ *Id.*

of unincorporated land that continued from her backyard into the desert.

She'd spent almost a decade in Pahrump^[191] without really knowing why. The heat could be unbearable. She had no family in Nevada. She loved going to movies, and the town of 30,000 didn't have a theater. It seemed to her like a place in the business of luring people — into the air-conditioned casinos downtown, into the legal brothels on the edge of the desert, into the new developments of cheap housing available for no money down — and in some ways she'd become stuck, too.

...

[She] didn't believe everything she read online, but she was also distrustful of mainstream fact-checkers and reported news. It sometimes felt to her like real facts had become indiscernible — that the truth was often somewhere in between. What she trusted most was her own ability to think critically and discern the truth, and increasingly her instincts aligned with the online community where she spent most of her time. It had been months since she'd gone to a movie. It had been almost a year since she'd made the hour-long trip to Las Vegas. Her number of likes and shares on Facebook increased each year until she was sometimes awakening to check her news feed in the middle of the night, liking and commenting on dozens of posts each day. She felt as if she was being let in on a series of dark revelations about the United States, and it was her responsibility to see and to share them.¹⁹²

The underlying problems suggested by this passage are so complex as to seem intractable (in addition to being unrelated to the information problems of fake news). And in any case, on some level, this consumer of fake news is getting what she wants. But should she? And how much is this due to the engineering of the Facebook experience?¹⁹³ Deeply contested questions of paternalist policy-making and what is normatively best for people follow closely behind.¹⁹⁴

While the fake news debate is particularly charged, similar issues arise with more mundane consumer issues. Perhaps AIs will tend, for example, to make recommendations that rest on past consumer experience at the expense of novelty,

¹⁹¹ Pahrump is in Nevada, about an hour west of Las Vegas.

¹⁹² Saslow, *supra* note 189.

¹⁹³ See, e.g., MCNAMEE, *supra* note 175, at 166 (describing how changes to Facebook's News Feed to favor algorithmic over human moderation “had the effect of promoting the primary elements of filter bubbles—family, friends, and Groups—at the expense of the content most likely to pierce a filter bubble, journalism”).

¹⁹⁴ To the extent these issues are resolvable, there remains the question whether the state can meaningfully regulate in this area given the First Amendment. See, e.g., Marc Jonathan Blitz, *Lies, Line Drawing, and (Deep) Fake News*, 71 OKLA. L. REV. 59, 116 (2018) (“The question raised by deep fakes and similar technology, then, is whether First Amendment law can leave government with room to protect the social foundations that allow individuals to serve as their own ‘watchmen for truth’ without simultaneously inviting officials to control and restrict how they play that role.”).

reflecting consumer inertia (in other words, consumers at the margin may, as a descriptive matter, prefer familiarity to change even when the change would be an improvement). What then?

Maybe the law will still want to drive consumers to consider new things in order to promote innovation and competition. Doing so would move the law in the direction of regulating consumer behavior. Resolving the issue would require development of a theory to explain when non-incumbent competitors should get a shot before the consumer.

Perhaps these issues should be addressed at a higher level of generality. Maybe the argument will be that if in the world to come AIs will have the power to guide us in a non-reflective state, then the law must require that people have more than one voice in their ears. Whatever our ability to find the appropriate vocabulary to debate these questions, we won't locate it in trademark law.

CONCLUSION

To some extent, trademark law is about authority. Who gets to define a mark and how? What kinds of uses may a markholder control? Some of these questions are easy. The owner of the COCA-COLA mark gets to determine what kind of soda bears the brand. One cannot redefine the mark by counterfeiting it. But one is free to comment on the trademarked product in ways that may affect mark meaning (e.g., "High-sugar products like COCA-COLA are a menace to public health").

Other issues are harder, and trademark law must constantly mediate between the claims of markholders and third parties making novel uses of trademark meanings. Many recent battles concerned the internet, and future technological developments will naturally continue to test trademark law. Imagining a hypothetical technological endpoint for digital assistants suggests ways that innovation may upset the place of trademark law in the consumer information ecosystem as we know it today. On the one hand, trademarks may lose importance as it becomes easier to turn purchasing decisions over to digital assistants or similar tools that look beyond the comparatively simple signals embodied in brand names. At the same time, digital technology is making it easier for different groups to attach different meanings to trademarks without experiencing the harms associated with trademark infringement.

These changes open the door to a more pluralist vision of trademark meaning, one that shifts the balance of power between trademark holders and the rest of us. Though that vision was not shared by the courts in the first wave of cases involving online trademarks, more recent precedent shows a greater receptivity.

But they also point to deeper issues regarding information in the digital era. Whatever the ultimate equilibrium between pluralist and authoritarian visions of trademark law, neither perspective has much to say to these challenges even though they might appear to touch on trademark law's domain. The thought experiment of this Article therefore supports those who see a larger gap in our laws that society has to decide whether and how to fill.

TRADEMARKS IN AN ALGORITHMIC WORLD

Christine Haight Farley*

Abstract: According to the sole normative foundation for trademark protection—“search costs” theory—trademarks transmit useful information to consumers, enabling an efficient marketplace. The marketplace, however, is in the midst of a fundamental change. Increasingly, retail is virtual, marketing is data-driven, and purchasing decisions are automated by AI. Predictive analytics are changing how consumers shop. Search costs theory no longer accurately describes the function of trademarks in this marketplace. Consumers now have numerous digital alternatives to trademarks that more efficiently provide them with increasingly accurate product information. Just as store shelves are disappearing from consumers’ retail experience, so are trademarks disappearing from their product search. Consumers may want to buy a product where the brand is the essential feature of the product such that the brand is the product, but they no longer need the assistance of a trademark to find the product.

By reflexively continuing to protect trademarks in the name of search costs theory, courts give only lip service to consumer interests without questioning whether trademarks are fulfilling any useful information function. In many cases, trademarks may actually misinform consumers by masking the identity of the producer or its distanced relationship with the trademark owner. Without having deliberately decided to do so, trademark law is now protecting “brands as property” without any supportive normative rationale. Removing the veil of search costs theory will enable courts to consider whether trademark protection is justified in particular cases.

- INTRODUCTION..... 1124
- I. THE NORMATIVE RATIONALE FOR TRADEMARKS 1130
 - A. Why Protect Trademarks?..... 1130
 - B. Trademarks Convey Source of Origin Information to Consumers..... 1131
 - C. Trademarks Reduce Consumers’ Search Costs 1134
- II. UNDERMINING TRADEMARK’S RATIONALE 1139
 - A. The Historically Contingent Value of Trademarks..... 1139
 - B. The Depreciation of Trademark’s Information Function.. 1145
 - C. Trademarks that Misinform Consumers 1148
- III. TRADEMARKS IN THE AGE OF ALGORITHMS 1154
 - A. The Tech Challenge to Trademark Theory 1154
 - B. The Disappearance of Trademarks from Search..... 1156

*Professor, American University Washington College of Law. Thanks to the participants in the Texas A&M Law School’s Conference on the Law & Economics of Trademarks, the University of Minnesota Law School’s Faculty Works in Progress lecture series, Boston University School of Law’s Faculty Works in Progress lecture series, and Stacey Dogan, Bill McGeeveran, Michael Meurer, Glynn Lunney, Andrew Ferguson, and David Simon in particular, for helpful feedback. I am also grateful for the excellent research assistance provided by Allison Clark, Jaclyn Bell, and Aasees Kaur.

1124	WASHINGTON LAW REVIEW	[Vol. 98:1123
	C. The Advent of Adtech.....	1162
	D. The Disappearance of Search and Ads	1170
IV.	REEVALUATING TRADEMARKS IN AN AI-DRIVEN MARKETPLACE	1173
	A. The Morphed Function of Trademarks	1176
	B. Shopping in the Metaverse.....	1179
	C. Trademarks Need a New Rationale	1181
	CONCLUSION	1185

INTRODUCTION

Wearing smart glasses, Jamie scans the room. Google instantly spots a desk chair and zooms in on it. Google knows Jamie has been sitting for long periods, has recently purchased an increased amount of ibuprofen, and has an upcoming chiropractic appointment. Google now captions the framed chair with the best place to purchase it, considering price, availability, shipping time and cost, and the rewards programs in which Jamie is enrolled. Jamie has come to rely on these smart recommendations without ever deciding to do so. This chair will be a good fit as it is based on Jamie's body dimensions, desk area (based on data derived from the robot vacuum—another Google suggestion), and style preferences. Jamie had not previously decided to purchase a chair, but was ready to then, having already been primed by news articles on back pain and the photos of stylish offices that populated Jamie's various social media feeds.

Back at home, the smart speaker senses Jamie's presence using ultrasound technology. Using facial and voice recognition, it identifies Jamie and listens in for future shopping suggestions even without the wake-up command. It announces that a grocery delivery, based on data from the smart fridge, will be delayed until tomorrow. Before Jamie can process this news, it has preemptively made a dinner reservation at a new restaurant based on Jamie's habits, calendar, and a contact's review. When Jamie awakes in the night and turns to the Kindle, Amazon records the time and book title for information about sleep patterns and interests, which will factor into future purchase suggestions.

Whether the above scenario sounds dystopian or convenient,¹ it illustrates the change that is underway in retail today. Consumers are

1. A Super Bowl commercial for Alexa captures both the convenience and intrusion of these technologies. When Scarlett Johansson wakes up in the morning and speaks to Colin Jost, their Echo announces that it is ordering extra strength mouthwash. See *Scarlett Johansson and Colin Jost in Super Bowl Ad*, CNN, <https://www.cnn.com/videos/media/2022/02/09/scarlett-johansson-colin-jost-super-bowl-ad-alexa-sot-france-nr-vpx.cnn/video/playlists/business-super-bowl/> (last visited Oct. 31, 2023).

increasingly relying on algorithms to make purchases.² These algorithms can accurately predict what a consumer will buy because the consumer's every move that can be surveilled has been collected and processed.³ Rather than consumers deciding to make a purchase or choosing what to purchase, they may receive curated purchase suggestions based on data analytics, or even presumptive purchases that can be cancelled should the consumer disagree with the machine. Even when a consumer exercises agency by digitally capturing an image of something they like on a screen or in the physical world, one click unleashes the machines to do the searching and sifting of information that results in an instantaneous purchase suggestion. In any event, the suggestion of what to buy and from whom is thoroughly researched, factoring in all of the consumer's preferences. Even the consumer's satisfaction with that purchase becomes an input for future suggestions.

Significantly, missing from this description of the new retail experience are trademarks. In these scenarios, consumers are not relying on trademarks to decide what to purchase. Instead, products are suggested to them based on a multitude of factors in which the trademark may not figure. When a consumer approves a purchase, trademarks may not be visible in the transaction. The machines may process the trademark, but the consumer may not. The fact that trademarks are increasingly less prevalent and less relevant to consumers in deciding what to purchase in this algorithm-mediated marketplace has so far escaped notice in the law or literature.

This Article critically re-examines trademark law's rationale by bringing surveillance capitalism to bear on trademark law. According to the sole theoretical justification for the law, trademarks function as concise identifiers of the source of products.⁴ By one glance, the trademark assures the consumer that the product is made by the same producer as a product they have previously experienced. Trademarks thus convey source information to consumers that assists them in selecting the product they desire. Articulated through the Chicago School of Economics as "search costs" theory, trademarks' social utility is that they enable an efficient marketplace by lowering consumer search costs.⁵ Since consumers have no practical capacity to do actual research on a product's origin, characteristics, or quality, trademarks function as a stand-in for

2. See *infra* section III.C.

3. See *infra* section III.C.

4. See *infra* section I.B.

5. See *infra* section II.C.

such research. Trademarks make shopping efficient by instantaneously transmitting source information to consumers.

The increasingly algorithmic world we inhabit reveals the faulty premise on which the rationale of trademark law rests. How consumers will make purchases in the future—and have started to already—stands in sharp contrast to trademark law’s account of how consumers behave in the marketplace. Trademark’s rationale thus depends on a descriptive account of the marketplace and how consumers behave in it. The marketplace, however, is not static, and that account is now anachronistic. The proliferation of algorithmic decision-making in the marketplace today exposes the historical contingency of the theoretical basis of trademark law and thus undercuts its rationale.

A look backward teaches that it was a similar dramatic change in the marketplace that spawned the modern theory of trademarks. Retail underwent a paradigm shift in the early twentieth century, which is when both consumer choice and the information theory of trademarks emerged.⁶ Before then, the problem that trademarks are said to address—efficiently finding and choosing which product to buy—was nonexistent. At that time, a product’s source of origin was of no concern to consumers because the source was usually anonymous. Instead, consumers relied on their local shopkeeper to choose which products to make available to them.⁷ The goodwill of consumers was an asset that the shopkeeper cultivated, not the producer.

We are now experiencing a second paradigm shift in retail where artificial intelligence (AI), algorithms, and advertising technology (adtech) play an increasingly central role. As with the pre-twentieth century marketplace, consumer choice is once again becoming mediated—this time by technology.⁸ As a result, like the pre-twentieth century marketplace, consumers will make purchases without relying on trademarks.

Even the marketplace of the twentieth century, on which the trademark information and search costs theories were based, failed to function as described. Gradual changes in the marketplace, such as increased assignments and licensing, subcontracted production, and complicated supply chains, undermined the trademark information theory. The relationship between the trademark owner and the origin of the product has become increasingly attenuated. These realities of the marketplace diminish the information function of trademarks and, in some cases,

6. See *infra* section III.A.

7. See *infra* section II.A.

8. See *infra* section II.A.

trademarks actually conceal source information that may be relevant to consumers' purchasing decisions.⁹ Search costs theory provides too simple an account of a much more complicated marketplace.

Thus far, neither scholarly critique nor the numerous inconsistent realities of the marketplace have been able to topple the persistence of search costs theory. However, with the advent of the AI-mediated marketplace we are now at an inflection point where search costs theory has met its justifying limit. This paradigm shift in the marketplace fundamentally undercuts the accepted rationale for protecting trademarks.

If we protect trademarks because they are efficient at conveying accurate information, it means that the cost of this protection outweighs the alternative, which would be less efficient or less accurate. The image of a consumer overwhelmed by a store shelf full of choices makes such a tradeoff seem necessary. Consumers today, however, possess alternative means to efficiently find, identify, and choose products. No longer must consumers rely solely on the informational shortcut provided by trademarks. Consequently, trademarks are vanishing from consumers' product searches.¹⁰ Consumers' increased technological capacity enables them to search for products, not trademarks. Trademarks no longer need to stand in for actual product research, which was formerly impractical.

Trademarks' justifying theory is based on consumer agency in purchasing decisions. But even that is disappearing as retail has become increasingly virtual, marketing ever more data-driven, and purchasing decisions driven by AI. Predictive analytics is changing consumer behavior. When it comes to making purchasing decisions, consumers are increasingly delegating the search for products, product selection, and whether to purchase out of convenience. Marketers today know vast amounts about us and have increasingly sophisticated means of using this knowledge to predict our purchasing choices.¹¹ Consumers, in turn, are recognizing that algorithms can make purchasing decisions faster and better than they can.¹²

This mediation of consumers' purchasing decisions fundamentally changes the flow of information. Rather than communicating with consumers through trademarks, producers will be increasingly interested in conveying product information to machines. Machine-to-machine communication offers accuracy and efficiency unparalleled by

9. See *infra* sections II.B–II.C.

10. See *infra* Part IV.

11. See *infra* section III.C.

12. See *infra* section III.C.

trademarks.¹³ In this marketplace, trademarks become superfluous—if not distracting—packets of information. Trademark law was supposed to foster a more decipherable marketplace for consumers. In surveillance capitalism, the marketplace is increasingly only decipherable by machines.¹⁴

Trademark law has thus far failed to acknowledge the fundamental disruptions to the marketplace brought about by digital technologies. Consumers used to watch ads; now, ads watch them.¹⁵ Now that the marketer has near-perfect information about the consumer, the consumer has little need for the informational value of the trademark.

The AI-driven marketplace undermines trademark law's sole rationale for protection. Today, and more so in the future, trademarks may offer no informational value or search efficiency and may even obfuscate such informational values. The correspondence between trademarks and the search costs rationale has been lost. At a minimum, if the information conveyance function can now be performed without the aid of trademarks, at least in some cases, the diminished value of trademarks should be accounted for in the law. To analyze the cost to competition of trademark protection, the benefit—the informational value of trademarks—must be reconsidered in light of the current availability of alternative means.

Trademarks are being protected even though they are not functioning as posited. Perversely, as the informational value of trademarks has diminished, the commercial value and level of legal protection of many trademarks has increased.¹⁶ These trademarks, or rather “brands,” have a different function in the marketplace. They are less packets of source information that travel with the product, and more the products themselves.¹⁷ The brand itself is the commodity exchanged. When consumers desire the brand as the product, they are not relying on the information function of trademarks. Brand value has supplanted trademark value, but trademark law has not yet acknowledged this changed marketplace function.

These new insights on the role of trademarks in the digital world provide the necessary foundation for considering the next big challenges for trademark law. Stuck in the twentieth century, trademark law is ill-prepared for the marketplace in the virtual world. In the metaverse, digital goods will be branded, but those brands will function to communicate

13. See *infra* section III.C. Although this technology offers accuracy, it is also prone to manipulation. See *infra* section III.C.

14. See *infra* Part III.

15. See *infra* section III.C.

16. See *infra* section IV.A.

17. See *infra* section IV.A.

status, image, aspirations, and associations from one consumer to another. Digital goods have no meaningful source of origin or hidden qualities or characteristics.¹⁸ They thus have no need for trademarks' informational function. Trademark law may well want to protect the actual functioning of brands in the metaverse, but it will need a new rationale to do so.

Trademark theory needs to acknowledge that trademarks are doing less work and doing it less often. Instead of protecting trademarks qua information, the law is protecting brands as property. Presently, brands enjoy robust protection because of unexamined application of search costs theory and not because of a considered determination that they are deserving of legal protection. This Article contends that the support for the economic welfare or social utility claim is now at its weakest point in the history of trademark law. It is time to recalibrate trademark law's normative foundations.

This Article proceeds as follows. Part I begins with an examination of the theoretical justification for trademark protection. The central idea is that trademarks assist consumers in the marketplace by imparting useful information in the most efficient manner. Accordingly, the theory that trademarks reduce consumer search costs now dominates trademark law. Part II observes that the search costs theory is fundamentally a descriptive account and, as such, is vulnerable to a changing marketplace. This Article briefly recounts the history of retail to show how trademark law developed as a response to a major change in the way consumers shopped. It then exposes the rationale for trademark protection as contingent on the particular way that marketplace functioned. Noting significant changes in the marketplace in the latter part of the twentieth century, Part II challenges the idea that trademarks invariably impart useful information to consumers. Trademark owners, who today are unlikely to be the producers of the products that bear their label, have become increasingly attenuated from the source of the products, and trademarks often mask complicated corporate relationships and untraceable supply chains. Part III contrasts search cost theory's imagined marketplace with today's AI-driven marketplace, where algorithms increasingly drive purchasing decisions. These data-driven purchases are highly efficient and are based on much more information than a trademark can impart. Consumers may not even be presented with trademarks in these transactions. Part IV argues that trademark law needs to acknowledge that trademarks are now doing less useful work for consumers and doing it less often. Rather than informing consumers about the source of the product, as search costs theory holds, many trademarks now aim to convey "brand attributes,"

18. See *infra* section IV.B.

such as associations, feelings, image, and status. The marketplace has shifted from trademarks to brands, but trademark law has failed to recognize this change because search costs theory cannot account for it. Instead of only protecting trademarks qua information, trademark law now protects brands as property. The uncritical acceptance of search costs theory has enabled the extension of trademark protection without even a question as to the social utility of such protection.

I. THE NORMATIVE RATIONALE FOR TRADEMARKS

A. *Why Protect Trademarks?*

The protection of trademarks is not easy to justify. As the United States Supreme Court explicitly observed in 1879 when evaluating the first federal trademark statute, trademarks, unlike copyrights and patents, do not involve creativity or inventiveness.¹⁹ Because trademarks were not identified by framers of the United States Constitution as the subject of federal legislation, again unlike copyrights and patents,²⁰ in order to receive federal protection, trademarks must affect interstate commerce.²¹ Thus, from the start, trademark law has been tied to the regulation of the marketplace without any more lofty policy goals. Early common law protection of trademarks was directed at routing out deceptive trade practices.²² Early trademark law consistently acknowledged that as the law was to facilitate the marketplace, the protection of trademarks should be limited so as not to harm competition.²³ Especially in light of trademark

19. Trade-Mark Cases, 100 U.S. 82, 93–94 (1879).

20. See U.S. CONST. art. I, § 8, cl. 8.

21. See *Trade-Mark Cases*, 100 U.S. at 94–97.

22. See 1 J. THOMAS MCCARTHY, MCCARTHY ON TRADEMARKS AND UNFAIR COMPETITION § 5:2 (5th ed. 2020) (“In 19th century cases, trademark infringement embodied many of the elements of fraud and deceit from which trademark protection developed.”); see also, e.g., *Pierce v. Guittard*, 8 P. 645, 647 (Cal. 1885) (“[I]t is a fraud on a person who has established a business for his goods and carries it on under a given name or with a particular mark, for some other person to assume the same name or mark, or the same with a slight alteration, in such a way as to induce persons to deal with him in the belief that they are dealing with a person who has given a reputation to that name or mark.”).

23. Beverly W. Pattishall, *Trade-Marks and the Monopoly Phobia*, 50 MICH. L. REV. 967, 968 (1952). Commentators, judges, and lawyers made frequent references to monopolies while discussing trademark policy. See, e.g., EDWARD CHAMBERLIN, THE THEORY OF MONOPOLISTIC COMPETITION 59–61 (1933) (discussing how firms in certain markets differentiate their products to gain a degree of market power); Pattishall, *supra*, at 968–69 (stating that many jurisdictions have adopted the “anti-monopoly” ideology when determining trademark protection). The legislative history of the Lanham Act reveals challenges brought by the United States Department of Justice opposing the Act’s creation of new and stronger rights by arguing that it was anticompetitive and might facilitate monopolies. See MCCARTHY, *supra* note 22, § 5:4.

law's interference with competition—as well as explicit accusations of monopoly rights²⁴—courts, treatise authors, and commentators sought to justify why protecting mere “labels” was socially valuable.²⁵ Two related justifications for the protection of trademarks emerged: the informational value of trademarks and their enhancement of market efficiency.

B. *Trademarks Convey Source of Origin Information to Consumers*

Today, the idea that trademarks provide valuable information to consumers, empowering them to efficiently search for the product they want to purchase, dominates trademark law. In fact, it is now the sole theoretical justification for the law.²⁶

According to the accepted justification, trademarks are consumers' helpers in the marketplace. The theory assumes that the source of the product is important to consumers, and that consumers seek out products from particular producers. It is accepted that consumers acquire knowledge about the product's source or producer based on their own prior experience with the products, the experience of others, and claims made by the producer.²⁷ This information then becomes a stand-in for actual research on the product's characteristics and quality, for which the consumer does not have capacity.²⁸ Because consumers are confronted with many choices within a product category and are not practically able to devote adequate time and consideration to most choices, they may make

24. See MCCARTHY, *supra* note 22, § 5:4.

25. See *infra* sections II.B–II.C.

26. Although Frank Schechter's 1927 *Harvard Law Review* article did articulate an alternate theory of trademarks, that underlying theory of trademarks never made any serious headway, even if trademark dilution was eventually added to the Lanham Act. See Robert G. Bone, *Schechter's Ideas in Historical Context and Dilution's Rocky Road*, 24 SANTA CLARA COMPUT. & HIGH TECH. L.J. 469, 472, 475, 505 (2008). As discussed *infra* notes 272–273, Schechter argued that the law's protection should respond to the way trademarks actually functioned in the marketplace, which was “to identify a product as satisfactory and thereby to stimulate further purchases by the consuming public.” See Frank I. Schechter, *The Rational Basis of Trademark Protection*, 40 HARV. L. REV. 813, 818 (1927) [hereinafter Schechter, *Rational Basis*]. The idea that trademarks have “selling power” has never been denied, but basing protection on that purpose has only seldom been acknowledged. See, e.g., *Mishawaka Rubber & Woolen Mfg. Co. v. S.S. Kresge Co.*, 316 U.S. 203, 205 (1942) (referring to trademark law as “the law's recognition of the psychological function of symbols”).

27. See William M. Landes & Richard A. Posner, *Trademark Law: An Economic Perspective*, 30 J.L. & ECON. 265, 269 (1987) [hereinafter Landes & Posner, *Economic Perspective*] (“[S]uppose a consumer has a favorable experience with brand X and wants to buy it again. Or suppose he wants to buy brand X because it has been recommended by a reliable source or because he has had a favorable experience with brand Y, another brand produced by the same producer. Rather than investigating the attributes of all goods to determine which one is brand X or is equivalent to X, the consumer may find it less costly to search by identifying the relevant trademark and purchasing the corresponding brand.”)

28. See *id.*

mistaken purchases. In this scenario, trademarks function as valuable packets of information that help consumers make informed purchasing decisions.²⁹ As stable bearers of immediate information, they enable an efficient marketplace. Trademarks are therefore worthy of protection because they ensure stability in the marketplace for both consumers and producers. Thus, trademark law protects any designation that consumers recognize as indicating the products' source of origin.³⁰

This idea that trademarks provide valuable information to consumers dates back to the earliest trademark cases. In its first trademark case in 1871, the United States Supreme Court stated that “[t]he office of a trademark is to point out distinctively the origin, or ownership of the article to which it is affixed; or, in other words, to give notice who was the producer.”³¹ Later, in 1916, in its influential *Hanover Star Milling Co. v. Metcalf*³² decision, the Court reiterated the utility of trademarks, stating that “[t]he primary and proper function of a trademark is to identify the origin or ownership of the article to which it is affixed.”³³ Just recently, in

29. See Robert G. Bone, *Hunting Goodwill: A History of the Concept of Goodwill in Trademark Law*, 86 B.U. L. REV. 547, 549 (2006) [hereinafter Bone, *Hunting Goodwill*] (“The core of trademark law, as it is understood today, is based on a model which I shall call the ‘information transmission model.’ This model views trademarks as devices for communicating information to the market and sees the goal of trademark law as preventing others from using similar marks to deceive or confuse consumers.”); Glynn S. Lunney, Jr., *Trademark Monopolies*, 48 EMORY L.J. 367, 369 (1999) (explaining the prevailing idea that “a trademark provides consumers with information that they need (and cannot otherwise readily obtain) in order to match their desires to particular products”); Stephen L. Carter, *The Trouble with Trademark*, 99 YALE L.J. 759, 759 (1990) (describing trademarks as “packets of information”); Stacey L. Dogan, *What Is Dilution, Anyway?*, 105 MICH. L. REV. FIRST IMPRESSIONS 103, 106 (2006) (“[T]rademark law has never aimed to provide exclusive rights in marks, but has focused on preserving informational clarity in the marketplace.”); Jake Linford, *Valuing Residual Goodwill After Trademark Forfeiture*, 93 NOTRE DAME L. REV. 811, 820 (2017) (“The law thus protects trademarks so that they can be used to convey accurate information”); William P. Kratzke, *Normative Economic Analysis of Trademark Law*, 21 MEM. ST. U. L. REV. 199, 216–17 (1991) (arguing the law should protect trademarks to facilitate the transmission of informational messages).

30. Lanham Act § 45, 15 U.S.C. § 1127; see also *B&B Hardware, Inc. v. Hargis Indus., Inc.*, 575 U.S. 138, 142 (2015) (“The principle underlying trademark protection is that distinctive marks . . . can help distinguish a particular artisan’s goods from those of others.”); *United Drug Co. v. Theodore Rectanus Co.*, 248 U.S. 90, 97 (1918) (noting a trademark “designat[e]s the goods as the product of a particular trader” and “protect[s] his good will against the sale of another’s product as his”); 1 ANNE GILSON L’ALONDE, *GILSON ON TRADEMARKS* § 1.03(1) (2023) (“[A] trademark . . . is accorded legal protection when it designates the source or origin of a product or service to distinguish that product or service from those of others, even if the consumer can’t name that source. . . . [Trademark law] makes it easier for consumers to choose the products and services they want.”).

31. *Canal Co. v. Clark*, 80 U.S. (13 Wall.) 311, 322 (1871).

32. 240 U.S. 403 (1916).

33. *Id.* at 412; see also Sidney A. Diamond, *The Historical Development of Trademarks*, 73 TRADEMARK REP. 222, 247 (1983) (“Trademarks are essential to the operation of a competitive system of free enterprise, for they are the only feasible means by which the consumer can select the

Jack Daniel's Properties, Inc. v. VIP Products LLC,³⁴ the Court gave emphasis to the notion that trademarks indicate source, designating this trademark's "'primary' function."³⁵ The Court acknowledged that "[t]rademarks can of course do other things: catch a consumer's eye, appeal to his fancies, and convey every manner of message."³⁶ But the Court emphasized that "whatever else it may do, a trademark is not a trademark unless it identifies a product's source (this is a Nike) and distinguishes that source from others (not any other sneaker brand)."³⁷ It further stated that the "primary mission" of the Trademark Act, known as the "Lanham Act," is plainly evident "[f]rom its definition of 'trademark' onward," as it "views marks as source identifiers—as things that function to 'indicate the source' of goods, and . . . 'distinguish' them from ones 'manufactured or sold by others.'"³⁸ In a case briefed and argued as a case about trademark fair use, the majority opinion used the word "source" forty-eight times in its relatively short, eleven-page opinion.³⁹

In these cases, the Supreme Court has adopted and standardized the notion that trademarks convey information about the origin, ownership, or identity of the producer of the product. This idea subsumes dependent concepts. The first is that the information the trademark conveys is that all products bearing the trademark come from a single source. The second is that because trademarks identify one seller's products, they help consumers distinguish these products from products sold by others. The third is that consumers can rely on trademarks to help them distinguish among different sellers' products because this information is stable. These concepts combine to produce the social utility of trademark information conveyance. At a minimum, with the ability to distinguish among products, trademarks offer consumers convenience in the marketplace. They allow consumers to sort. It is sometimes stated that the distinguishing function of trademarks is essential, rather than just a convenience.⁴⁰ Whether or not consumers want to sort by trademark

particular variety of product that he wishes to buy from among the multitude of choices that manufacturers now make available to satisfy individual tastes and individual preferences among the purchasing public.").

34. 599 U.S. ___, 143 S. Ct. 1578 (2023).

35. *Id.* at 1583.

36. *Id.*

37. *Id.*

38. *Id.* at 1589 (quoting 15 U.S.C. § 1127).

39. *Id.* at 1578.

40. *See, e.g., Lunney, supra* note 29, at 369 ("By identifying the source of goods, a trademark provides consumers with information that they need (and cannot otherwise readily obtain) in order to match their desires to particular products."); S. REP. NO. 79-1333, at 4 (1946) ("Trade-

depends on whether consumers have preexisting information about a producer.⁴¹ There is a certain tautology in this theory: trademarks function by conveying source information, so long as the law protects them from being copied; consumers will rely on trademarks for this source information; consumers will even come to invest these designations with their product research; consumers can then conveniently match products in the marketplace with their desires. This social utility deserves legal protection.

C. Trademarks Reduce Consumers' Search Costs

The information function of trademarks adopted by the Supreme Court was further shaped by the Chicago School of Economics. William Landes and Richard Posner set their law and economics lens on trademark law in 1987 and offered a positive economic justification for trademarks. They rearticulated the prevailing understanding of how consumers rely on trademarks in the marketplace as a "search costs theory."⁴² They argued that the social utility of trademarks is that they enable an efficient marketplace by lowering consumer search costs.⁴³ Trademarks perform this function by instantaneously transmitting source information to consumers to assist them to select the product they desire.⁴⁴ Applying this theory, Judge Posner explained that

[t]he fundamental purpose of a trademark is to reduce consumer search costs by providing a concise and unequivocal identifier of

marks . . . make possible a choice between competing articles by enabling the buyer to distinguish one from the other.").

41. See, e.g., S. REP. NO. 79-1333, at 3 (1946) ("[One] purpose underlying any trade-mark statute . . . is to protect the public so it may be confident that, in purchasing a product bearing a particular trade-mark which it favorably knows, it will get the product which it asks for and wants to get.").

42. Landes & Posner, *Economic Perspective*, *supra* note 27, at 269; William M. Landes & Richard A. Posner, *The Economics of Trademark Law*, 78 TRADEMARK REP. 267, 272 (1988) [hereinafter Landes & Posner, *Economics of Trademark*]; see also Nicholas S. Economides, *The Economics of Trademarks*, 78 TRADEMARK REP. 523, 526-27 (1988) ("[S]ellers have much better information as to the unobservable features of a commodity for sale than the buyers. . . . Unobservable features, valued by the consumer, may be crucial determinants of the total value of the good. . . . [I]f there is a way to identify the unobservable qualities, the consumer's choice becomes clear. . . . The economic role of the trademark is to help the consumer identify the unobservable features of the trademarked product. This information is not provided to the consumer in an analytic form, such as an indication of size or a listing of ingredients, but rather in summary form, through a symbol which the consumer identifies with a specific combination of features."); S. REP. NO. 79-1333, at 4 (1946) (noting that trademarks "make possible a choice between competing articles by enabling the buyer to distinguish one from the other").

43. Landes & Posner, *Economic Perspective*, *supra* note 27, at 269.

44. See *id.*

the particular source of particular goods. The consumer who knows at a glance whose brand he is being asked to buy knows whom to hold responsible if the brand disappoints and whose product to buy in the future if the brand pleases.⁴⁵

Landes and Posner use the example of the consumer who spots the trademark Crest on a store shelf.⁴⁶ That consumer instantly knows the source of the toothpaste and can select it from the many others on offer. The source of that toothpaste is the Proctor and Gamble Company,⁴⁷ but the consumer need not know the actual identity of the producer, only that Crest means it comes from a single source. This is the information that enables the consumer to distinguish this product from other toothpastes, thus reducing their search costs.

The significant contribution made by the Chicago School to the normative theory of trademarks is simply the reframing of the information theory of trademarks as an efficiency concern. The focus is on the cost to consumers in terms of time and effort searching within a product category for a selection that best meets their needs and desires. According to Landes and Posner, without the aid of the trademark, consumers would be forced to spend time researching each product choice.⁴⁸

45. Ty Inc. v. Perryman, 306 F.3d 509, 510 (7th Cir. 2002).

46. Landes & Posner, *Economics of Trademark*, *supra* note 42, at 278.

47. *Id.*

48. *See id.* at 277.

This economic rationale has been repeated and rearticulated by courts,⁴⁹ treatise authors,⁵⁰ and scholars⁵¹ ever since. The Supreme Court

49. See, e.g., *Scandia Down Corp. v. Euroquilt, Inc.*, 772 F.2d 1423, 1429 (7th Cir. 1985) (noting that trademarks “convey valuable information to consumers at lower costs” because they “reduce the costs consumers incur in searching for what they desire, and the lower the costs of search the more competitive the market”); *W.T. Rogers Co. v. Keene*, 778 F.2d 334, 338 (7th Cir. 1985) (“The purpose [of trademarks] is to reduce the cost of information to consumers by making it easy for them to identify the products or producers with which they have had either good experiences, so that they want to keep buying the product (or buying from the producer), or bad experiences, so that they want to avoid the product or the producer in the future.”); *Brennan’s, Inc. v. Brennan’s Rest., L.L.C.*, 360 F.3d 125, 132 (2d Cir. 2004) (“[I]f the mark is not recognized by the relevant consumer group, a similar mark will not deceive those consumers and thereby increase search costs.”); *New Kids on the Block v. News Am. Publ’g., Inc.*, 971 F.2d 302, 305 n.2 (9th Cir. 1992) (“[T]rademarks reduce consumer search costs by informing people that trademarked products come from the same source.”); *Kohler Co. v. Moen Inc.*, 12 F.3d 632, 643 (7th Cir. 1993) (“Innovation in product design and marketing for the purpose of enhancing producer identity reduces the costs to consumers of informing themselves about the product source so that they can . . . continue purchasing the products from particular producers”); *Kraft Foods Grp. Brands LLC v. Cracker Barrel Old Country Store, Inc.*, 735 F.3d 735, 739 (7th Cir. 2013) (“A trademark’s value is the saving in search costs made possible by the information that the trademark conveys about the quality of the trademark owner’s brand.”); *Bretford Mfg., Inc. v. Smith Sys. Mfg. Corp.*, 419 F.3d 576, 579 (7th Cir. 2005) (“Trademark law is designed to reduce the costs customers incur in learning who makes the product”); *Bos. Duck Tours, LP v. Super Duck Tours, LLC*, 531 F.3d 1, 14 (1st Cir. 2008) (noting that “the primary justifications for protecting trademarks” are “to aid competition and lower consumers’ search costs”); *Union Nat’l Bank Tex., Laredo v. Union Nat’l Bank Tex., Austin*, 909 F.2d 839, 844 (5th Cir. 1990) (“[T]rademarks are ‘distinguishing’ features which lower consumer search costs and encourage higher quality production by discouraging free-riders.”).

50. See, e.g., *MCCARTHY*, *supra* note 22, § 2:3 (“[A] trademark is a symbol that allows a purchaser to identify goods or services that have been satisfactory in the past [T]hey reduce the customer’s costs of shopping and making purchasing decisions.”); *GILSON LALONDE*, *supra* note 30, § 1.03(6) (“Trademarks enable consumers to purchase the specific goods they intend to purchase, reduce search costs, and help avoid deception in the marketplace.”).

51. See, e.g., *Carter*, *supra* note 29, at 759 (“[Trademarks] lower consumer search costs, thus promoting the efficient functioning of the market.”); Stacey L. Dogan & Mark A. Lemley, *Trademarks and Consumer Search Costs on the Internet*, 41 HOUS. L. REV. 777, 787–88 (2004) [hereinafter Dogan & Lemley, *Trademarks and Consumer Search*] (“Consumers benefit because they don’t have to do exhaustive research or even spend extra time looking at labels before making a purchase”); Robert G. Bone, *Taking the Confusion Out of “Likelihood of Confusion”*: *Toward a More Sensible Approach to Trademark Infringement*, 106 NW. U. L. REV. 1307, 1311–12 (2012) [hereinafter Bone, *Trademark Infringement*] (“[T]rademarks reduce the costs to consumers of searching for product information.”); Bone, *Hunting Goodwill*, *supra* note 29, at 555 (“[P]rotecting the source identification and information transmission function of marks . . . helps to reduce consumer search costs.”); William McGeeveran, *Rethinking Trademark Fair Use*, 94 IOWA L. REV. 49, 54 (2008) (noting that trademarks “reduce inefficient search costs”); Clarisa Long, *Dilution*, 106 COLUM. L. REV. 1029, 1056–59 (2006) (“[T]rademarks serve as source identifiers for consumers and thereby reduce consumer search costs.”); Maureen A. O’Rourke, *Defining the Limits of Free-Riding in Cyberspace: Trademark Liability for Metatagging*, 33 GONZ. L. REV. 277, 307 n.114 (1997) (“One of the goals of trademark law is to reduce consumer search costs.”); Margreth Barrett, *Internet Trademark Suits and the Demise of “Trademark Use,”* 39 U.C. DAVIS L. REV. 371, 376–78 (2006) (“[P]reventing misleading uses of marks that may confuse consumers about the source, sponsorship, or affiliation of the products or services they buy . . . reduces consumer search costs”); Ariel

has adopted this refinement of trademark theory, stating in *Qualitex Co. v. Jacobson Products Co.*⁵² that “trademark law . . . ‘reduce[s] the customer’s costs of shopping and making purchasing decisions.’”⁵³ The Court elaborated that a trademark “quickly and easily assures a potential customer that *this* item—the item with this mark—is made by the same producer as other similarly marked items that he or she liked (or disliked) in the past.”⁵⁴

Most commentators accept the idea that trademarks “facilitate the transmission of accurate information to the market,”⁵⁵ whether they support or critique the extent of the law’s protections of trademarks.⁵⁶ The search costs theory has become the blindingly dominant theoretical justification of trademark law.⁵⁷ In fact, no alternative account of

Katz, *Beyond Search Costs: The Linguistic and Trust Functions of Trademarks*, 2010 BYU L. REV. 1555, 1557 (2010) (“[T]rademarks are socially beneficial because they reduce consumer search costs.”); Michael Grynberg, *The Road Not Taken: Initial Interest Confusion, Consumer Search Costs, and the Challenge of the Internet*, 28 SEATTLE U. L. REV. 97, 97–99 (2004) (proposing a new framework, focused on consumer search costs, for analyzing initial interest claims); Lunney, *supra* note 29, at 432 (“Trademarks can, thus, reduce the searching costs involved in identifying the desired product.”); David W. Barnes, *Trademark Externalities*, 10 YALE J.L. & TECH. 1, 35–36 (2007) (“Because reducing consumers’ search costs is one of the goals of trademark law, preventing this increase in search costs is a benefit of recognizing exclusive rights.”); Linford, *supra* note 29, at 819 (“Trademark law is thus efficient to the extent it reduces consumer search costs . . .”).

52. 514 U.S. 159 (1995).

53. *Id.* at 163–64 (alteration in original) (quoting 1 J. THOMAS MCCARTHY, MCCARTHY ON TRADEMARKS AND UNFAIR COMPETITION § 2.01(2) (3d ed. 1994)).

54. *Id.* at 164.

55. *E.g.*, Bone, *Hunting Goodwill*, *supra* note 29, at 548; *id.* (noting that trademark law’s “core mission, as it is understood today, is to facilitate the transmission of accurate information to the market”).

56. Some scholars whose work argues for checks on the expansion of trademark rights support the search costs theory because the alternative—a property theory of trademarks—would likely produce further expansion of rights that would not be in the consumers’ interest. *See, e.g.*, Dogan & Lemley, *Trademarks and Consumer Search*, *supra* note 51, at 782 (“We argue that, cumulatively, these two axes of trademark expansion pose a grave danger to the law’s information-facilitating goals.”); Bone, *Trademark Infringement*, *supra* note 51, at 1378 (proposing “restructuring the infringement test so that it more closely fits the policies trademark law should serve”); Stacey L. Dogan & Mark A. Lemley, *A Search-Costs Theory of Limiting Doctrines in Trademark Law*, 97 TRADEMARK REP. 1223, 1223 (2007) [hereinafter Dogan & Lemley, *Search-Costs*] (“While the search costs theory provides a compelling argument for trademark rights, it also compels an equally important—but often overlooked—set of principles for defining and limiting those rights.”).

57. *See* Mark A. Lemley, *The Modern Lanham Act and the Death of Common Sense*, 108 YALE L.J. 1687, 1690 (1999) (citing Nicholas Economides, George J. Stigler, and others who fit within the Chicago School approach and stating that trademarks “communicate useful information to consumers, and thereby reduce consumer search costs”); Carter, *supra* note 29, at 762 (“The economic argument for protecting marks is straightforward and quite forceful. The principal benefit of trademark protection is that it lowers consumer search costs.” (footnote omitted)).

trademark policy even exists.⁵⁸ The theory has so firmly taken hold that legal theorists have been unable uproot it.⁵⁹

The increasingly algorithmic world we inhabit exposes trademark law's rationale as inadequate. The proliferation of algorithmic decision-making undercuts the search costs theory by revealing its faulty premise. The search costs theory depends on a descriptive account of the marketplace and how consumers behave in it. This account hinges on a series of descriptive claims about how a consumer makes purchasing decisions. First, consumers are often motivated to make purchases based on previous positive experiences with a particular producer and their interest in repeating that experience. Second, consumers have limited ability to research products or their producers and rely on trademarks as a proxy for such information. Third, consumers are often unwary and easily distracted. Fourth, bad actors sometimes take advantage of this situation. Based on the consistent repetition of these descriptive claims in trademark cases alone, one could form a distinct mental picture of the marketplace: the consumer stands in a store before shelves of products from many different producers scanning for the particular trademark that has previously been satisfactory. This picture, however, stands in sharp contrast to the way that Jamie made purchases as described at the beginning of the Article.

58. Barton Beebe, *The Semiotic Analysis of Trademark Law*, 51 UCLA L. REV. 621, 624 (2004); see also Mark P. McKenna, *A Consumer Decision-Making Theory of Trademark Law*, 98 VA. L. REV. 67, 75 (2012) [hereinafter McKenna, *Consumer Decision-Making Theory*] (“[I]t would be nearly impossible to overstate the extent to which the search costs theory now dominates as the theoretical justification of trademark law.” (emphasis in original)); Barton Beebe, *Search and Persuasion in Trademark Law*, 103 MICH. L. REV. 2020, 2069 (2005) [hereinafter Beebe, *Search and Persuasion*] (“The received orthodoxy [is] that trademark law exists to minimize our search costs . . .”); Deven R. Desai, *The Chicago School Trap in Trademark: The Co-Evolution of Corporate, Antitrust, and Trademark Law*, 37 CARDOZO L. REV. 551, 557 (2015) (“[T]he core logic of the Chicago School has not only taken over and now drives corporate and antitrust law but also drives modern trademark law.”).

59. See, e.g., Jeremy N. Sheff, *Marks, Morals, and Markets*, 65 STAN. L. REV. 761, 762 (2013) (arguing that search costs theory cannot account for much of trademark law's recent development); Jeremy N. Sheff, *Biasing Brands*, 32 CARDOZO L. REV. 1245, 1254 (2011) [hereinafter Sheff, *Biasing Brands*] (“[T]he search-costs model is descriptively inaccurate, or at least incomplete . . .”); Desai, *supra* note 58, at 619 (“[T]he search costs and competition rationale behind current trademark law and theory flows from and serves the Chicago School vision of autonomous firms wielding all resources to maximize wealth as firms see fit.”); McKenna, *Consumer Decision-Making Theory*, *supra* note 58, at 72 (“Rather than targeting search costs or confusion in and of themselves, trademark law should instead focus on deceptive practices that interfere with consumers' purchasing decisions.”); Irina D. Manta, *Hedonic Trademarks*, 74 OHIO ST. L.J. 241, 247 (2013) (“[T]he search-costs theory is incomplete, as scholars have pointed out, and does not comport with some of the existing trademark doctrines.”); see also Ann Bartow, *Likelihood of Confusion*, 41 SAN DIEGO L. REV. 721, 816 (2004) (“The educational and informational assistance trademarks provide to consumers is often overstated . . .”).

The changing technological landscape undermines the dominant rationale of trademark protection by exposing the historical contingency of its descriptive account. New technology demonstrates that the marketplace that served as the basis of the justification for protecting trademarks is rapidly becoming obsolete.

II. UNDERMINING TRADEMARK'S RATIONALE

A. *The Historically Contingent Value of Trademarks*

The digital marketplace that consumers encounter today has little connection to search costs theory. Contrasting the digital marketplace with search costs theory exposes the contingency of this theory on a particular conception of the marketplace. In a Seventh Circuit opinion that is thought to be the source of Landes and Posner's search costs theory,⁶⁰ Judge Easterbrook stated that "[t]rademarks help consumers to select goods."⁶¹ Trademark assistance consists of "identifying the source of the goods" and thus "convey[ing] valuable information to consumers."⁶² Significantly for Easterbrook and the Chicago School, this assistance "lower[s] the costs of search."⁶³ Easterbrook also attached trademark infringement to the search costs theory, explaining that "[c]onfusingly similar marks make consumers' task in searching for products harder."⁶⁴

Considering that this is the entire theoretical basis upon which the legal protection of trademarks rests, it is surprisingly limited. The theory of trademarks or information regulation it offers is inseparable from its description of how consumers interact with trademarks in the contemporary marketplace. As a descriptive account of the contemporary

60. See Jerre B. Swann, *The Evolution of Trademark Economics—from the Harvard School to the Chicago School to WIPO 2013—as Shepherded by INTA and The Trademark Reporter*, 104 TRADEMARK REP. 1132, 1135 (2014). Shortly after the *Scandia Down Corp. v. Euroquilt, Inc.*, 722 F.2d 1423 (1985), opinion by Judge Easterbrook, Judge Posner wrote the majority opinion in a case and rearticulated the search costs theory. See *W.T. Rogers Co. v. Keene*, 778 F.2d 334, 338 (7th Cir. 1985). While the *Scandia Down* opinion cites no authority for the search costs theory—it is stated as if a common understanding—the *W.T. Rogers* opinion cites both *Scandia Down* and a student note from 1984. See *id.* (first citing Peter E. Mims, Note, *Promotional Goods and the Functionality Doctrine: An Economic Model of Trademarks*, 63 TEX. L. REV. 639, 656–62 (1984); and then citing *Scandia Down*, 772 F.2d at 1429–30). Mims' note relies on Ralph H. Folsom & Larry L. Teply, *Trademarked Generic Words*, 89 YALE L.J. 1323 (1980). Mims, *supra*, at 641 n.7, 658 n.102. The Folsom and Teply article adopts the economic concept of search costs from George J. Stigler, *The Economics of Information*, 69 J. POL. ECON. 213 (1961), and applies it to trademark law, apparently for the first time. Folsom & Teply, *supra*, at 1335 nn.67–68 & 70.

61. *Scandia Down*, 772 F.2d at 1429.

62. *Id.*

63. *Id.*

64. *Id.* at 1430.

marketplace, it is vulnerable. First, Seventh Circuit judges and University of Chicago Law School professors are probably ill-suited to offer such a descriptive account. Search costs theory may, in fact, be based more on intuition or imagination than observation or empirical evidence.⁶⁵ Often when describing search costs theory, the example that proponents proffer to illustrate the theory is the ordinary consumer in a supermarket searching the shelves.⁶⁶ In the 1980s, the person ordinarily in the supermarket would most likely have been a housewife. It is possible, then, that the descriptive account is built on an outmoded stereotype.⁶⁷

Second, and more significant, search costs theory is vulnerable as a descriptive account due to its historical contingency. Even if it accurately described the 1980s marketplace, search costs theory should not be mistaken for a universal principle. Just as the above description of the digitally enabled consumer does not match up with search costs theory's account, a review of retail's history in America reveals it does not accurately describe the purchasing process of ordinary consumers in earlier marketplaces. Such time-stamped snapshots of the marketplace reveal search costs theory's limits.

Trademarks emerged to respond to a specific marketplace. Trademark law is based on a particular conception of that marketplace and how consumers function within it. But that marketplace is time-bound,

65. It could also be argued that the premises that consumers want to spend less time shopping or that shopping is just about buying something you need, do not accurately describe all consumers. Today, retail firms offer consumers various in-store and online experiences, such as AR, because they have found that consumers value spending time in this way. See Bobby Marhamat, *Five In-Store Customer Experiences Every Retailer Should Provide*, FAST CO. (Jan. 14, 2022), <https://www.fastcompany.com/90712962/five-in-store-customer-experiences-every-retailer-should-provide> (last visited Oct. 13, 2023). In the new experience economy, consumers buy so that they may shop more. Cf. *id.*

66. See, e.g., MCCARTHY, *supra* note 22, § 2:5 (“No busy working person in a developed society has hours to spend agonizing anew over every single purchase at the supermarket or elsewhere.”); Lunney, *supra* note 29, at 393 (“Consumer desire for, and the efficiency advantages associated with, trademark protection that will enable consumers to identify easily the particular soda they want from a shelf full of sodas seem reasonably clear.”); Economides, *supra* note 42, at 527–28 (“A typical example of such a product is a bottle of diet COKE, the cola beverage. Information on the bottle and label give little indication of the taste. The trademark identifies the product. A consumer is typically offered a free introductory bottle, or buys the first bottle to sample it. From his experience he is then able to decide rationally as an informed consumer about his future choices between diet COKE and all other goods.”).

67. In the context of trademark infringement litigation, Ann Bartow has similarly observed judges' tendency “to inappropriately rely on personal intuition and subjective, internalized stereotypes,” Bartow, *supra* note 59, at 723, and the impact in the law of such gender stereotyping. See *id.* at 776–92; see also *id.* at 777 (“[O]ne notices that many of the most demeaning conceptions of consumers, and of their intelligence and discriminatory powers (or lack thereof), tend to be articulated in trademark cases in which the underlying goods and services are primarily designed for, marketed to, or purchased by women.”).

resembling neither today's marketplace nor the marketplace of the nineteenth and early twentieth century. The first reported federal trademark case was decided in 1844.⁶⁸ Just as the marketplace has changed between 1987 and now, it certainly underwent significant changes between 1844 and 1987.⁶⁹ The changes in the law over that period reflect those changes in the marketplace.⁷⁰ However, since Landes and Posner released their article in 1987, trademark law's conception of the marketplace has been shockingly stable.⁷¹

Search costs theory is a refinement of a preexisting idea about trademarks developed in common law that trademarks provide consumers with relevant information about the products. That idea was developed in response to a perception about how consumers navigated the marketplace that then existed.⁷² The rationale of trademark protection has depended on the particular characteristics of the marketplace. Should those characteristics fundamentally change, the rationale for trademark protection should be adjusted accordingly.

Although trademark rights developed in the courts beginning in the 1840s in the United States, they were on shaky ground until the first decades of the twentieth century when they began to find their bearings. Consumers had little need for trademarks until then.⁷³ In the eighteenth and nineteenth centuries, consumers did not behave as search costs theory suggests.

The first retail establishments in the early nineteenth-century United States were local general stores run by shopkeeper-owners.⁷⁴ During this time, consumers typically made all of their purchases at their local general

68. See FRANK I. SCHECHTER, *THE HISTORICAL FOUNDATIONS OF THE LAW RELATING TO TRADEMARKS* 134 (1925) (citing *Taylor v. Carpenter*, 23 F. Cas. 742 (C.C.D. Mass. 1844) (No. 13,784)). The first state trademark case was decided in 1837. *Id.* (citing *Thompson v. Winchester*, 36 Mass. (19 Pick.) 214 (Mass. 1837)). See generally *id.* ("Up to 1870 only sixty-two trade-mark cases in all were decided by American courts.").

69. See SUSAN STRASSER, *SATISFACTION GUARANTEED: THE MAKING OF THE AMERICAN MASS MARKET* 19 (1989).

70. See *supra* section II.B.

71. See, e.g., *Peaceable Planet, Inc. v. Ty, Inc.*, 362 F.3d 986, 992 (7th Cir. 2004) (noting the importance and context of identifiability to the consumer within trademark); see also Desai, *supra* note 58, at 555 (discussing the influence of the Chicago School of Law and Economics on antitrust and corporate law).

72. Landes & Posner, *Economic Perspective*, *supra* note 27, at 269.

73. See Fred Mitchell Jones, *Retail Stores in the United States 1800–1860*, 1 J. MKTG. 134, 134 (1936).

74. See *id.*; see also PAMELA WALKER LAIRD, *ADVERTISING PROGRESS: AMERICAN BUSINESS AND THE RISE OF CONSUMER MARKETING* 15 (2020).

store.⁷⁵ The shopkeeper sourced the merchandise from the traveling salesperson who worked for large wholesale houses and manufacturers from the large cities.⁷⁶ In this way, shopkeepers “curated” the consumer products available to consumers. Consumers had little to no choice within a product category.⁷⁷ The problem of efficiently choosing *which* product to buy was nonexistent.

Consumers knew their local shopkeepers and had to trust them. Shopkeeper’s goodwill was generally dependent on the quality of the products that they sold.⁷⁸ Customers mostly bought on credit, which also produced strong loyalty.⁷⁹ It would not have occurred to consumers to care about the source of origin of the products; their satisfaction with the quality of products would affect the shopkeeper’s goodwill, not the producer’s.⁸⁰

The source of the products on offer was anonymous. At this time, it would not have been clear to consumers whether products came from a single source or not. Most of the products sold at these local shops were unbranded.⁸¹ Some products were sold in packages stamped with the shopkeeper’s name as “sole agent.”⁸² The source of origin of the products was either perfectly obvious because it was local, or immaterial because the interface with the shopkeeper was the only source of concern to consumers.⁸³ Just as AI is today, the shopkeeper was an intermediary between the consumer and their purchasing decision and succeeded in this role to the extent that they satisfied their customers’ desires.

75. See STEPHEN H. PROVOST, *THE GREAT AMERICAN SHOPPING EXPERIENCE* 1–2 (2021) (noting that most towns only had one general store and that general stores supplied all the necessities); Mitchell Jones, *supra* note 73, at 134 (“The inhabitants of the thinly settled regions had no other source of supply of importance [than general stores], with the exception of the peddler, for those goods which they required and could not produce by their own efforts.”).

76. SUSAN V. SPELLMAN, *CORNERING THE MARKET: INDEPENDENT GROCERS AND INNOVATION IN AMERICAN SMALL BUSINESS* 83 (2016). Traveling salesmen were called “drummers” because they drummed up business for wholesalers. *Id.*; see also STRASSER, *supra* note 69, at 19.

77. QUENTIN R. SKRABEC, JR., *THE 100 MOST SIGNIFICANT EVENTS IN AMERICAN BUSINESS* 119 (2012) (describing how Heinz used branding to change the way consumers interacted at the store by differentiating his products from “unlabeled competitors” so they ask for the Heinz brand name).

78. See Bone, *Hunting Goodwill*, *supra* note 29, at 575.

79. BENJAMIN LORR, *THE SECRET LIFE OF GROCERIES* 25 (2021).

80. See Bone, *Hunting Goodwill*, *supra* note 29, at 575.

81. SKRABEC, *supra* note 77, at 119 (“Most food products in the 1800s were sold as unbranded commodities in barrels, jars, and bags. Manufacturers sold to wholesalers who then distributed barrels of pickles, ketchup, flour, and so on to local stores.”).

82. JAMES D. NORRIS, *ADVERTISING AND THE TRANSFORMATION OF AMERICAN SOCIETY, 1865–1920*, at 14 (1990).

83. See Bone, *Hunting Goodwill*, *supra* note 29, at 575. A 1920 book argued against expanded trademark protection because it diminished the role of merchants as expert advisers to the consuming public. SIDNEY WEBB & BEATRICE WEBB, *INDUSTRIAL DEMOCRACY* 683–84 (1920).

Also analogous to online shopping today, consumers in this period did not interact with products until after they were purchased. Consumers often passed their shopping lists to clerks who collected the products for the consumer.⁸⁴ Inaccessible to the consumer, products were typically stored on shelves shielded by a counter and clerk.⁸⁵

Marketplaces began to shift from local to national in the early part of the twentieth century with the technology of mass production.⁸⁶ One innovation that changed retail was the development of packaging materials. Pre-cut cardboard boxes, invented in 1890, tin cans, paper bags, and card stock packaging allowed manufacturers to sell more products in packages.⁸⁷ Branding followed directly from this development. If products were sold in a box, then a name had to be put on the box. Products previously sold in bulk and unlabeled suddenly became branded.⁸⁸

Advertising followed the emergence of national markets.⁸⁹ Producers used advertising to convince consumers to abandon homemade and bulk goods and in exchange to buy factory-produced branded merchandise.⁹⁰ Because these goods often could not compete on cost or quality, advertisers resorted to appeals that played on emotions and status consciousness.⁹¹ If producers' direct appeals to consumers in radio advertisements were successful than consumers might ask for these products in their local stores. In this way, early radio ads could function as brand name prescription drug ads do today when they tell consumers to ask their doctor for a particular brand. Advertisers relied on customers to request items from retailers who did not carry the product.⁹² This was the first instance of trademarks serving as a proxy for buyer security. The message was clear: trust trademarks; seek them out. Trademark practitioner and historian Frank Schechter quoted author H.G. Wells on the advent of this shift: "[F]irms reach[ed] their hands over the retail

84. See LORR, *supra* note 79, at 23–25.

85. *Id.*

86. See STRASSER, *supra* note 69, at 6–7, 15–17.

87. LORR, *supra* note 79, at 26.

88. One example is National Biscuit Company, the manufacturer of Ritz Crackers. *Id.* at 27.

89. Desai, *supra* note 58, at 557 (“[A]dvertising and branding practices helped firms move beyond local retailers and reach consumers directly.”); see also Edward S. Rogers, *Some Historical Matter Concerning Trade-Marks*, 62 TRADEMARK REP. 239, 253 (“[T]he increased use of trade marks, brands and other identifying indicia in present day business . . . is directly traceable to the more extended distribution of products and this in turn is either the cause or the effect of modern advertising.”).

90. See, e.g., SKRABEC, *supra* note 81, at 119 (noting that Heinz used advertising to convince consumers to buy Heinz-branded products instead of bulk goods).

91. See Bone, *Hunting Goodwill*, *supra* note 29, at 580–81.

92. See STRASSER, *supra* note 69, at 21.

tradesman's shoulder, so to speak, and offer[ed] their goods in their own name to the customer."⁹³

Shopkeepers, however, remained an intermediary between consumers and national brands for some time. Advertising enabled consumers to recall a brand name, but they still had to ask the store clerk to retrieve products from the walls of shelving behind the counter.⁹⁴ Therefore, shopkeepers could substitute other manufacturers' products if they did not carry that brand or preferred to sell another.⁹⁵

Consumers were not able to examine products without the assistance of a shopkeeper until the 1916 invention of the self-service store.⁹⁶ Modeled on the invention of the cafeteria, consumers, supplied with a basket, followed a set pathway past every packaged product for sale.⁹⁷ This was a significant marketplace shift, as the absence of the intermediary shopkeeper meant that the items now had to speak for themselves.

The nature of shopping fundamentally changed at this moment. Now shopping was an opportunity to exercise personal preference. Selection of the right product was presented as a weighty decision. Consumers were encouraged to demonstrate their skill in making purchasing choices that saved money and provided quality products for their families.⁹⁸

Other retail innovations bolstered consumer choice. National catalog companies, such as Montgomery Ward and Sears and Roebuck, brought unprecedented choices to every household in the United States.⁹⁹ Supermarkets, which came onto the scene in 1930, stunned consumers with the sheer volume of products for consideration in an unmediated manner.¹⁰⁰ The shopping cart was invented in 1937 because, for the first time, consumers were buying more products than they could carry.¹⁰¹ Supermarkets quickly proliferated. In just twenty years, nearly every

93. Schechter, *Rational Basis*, *supra* note 26, at 818 n.21 (emphasis omitted).

94. STRASSER, *supra* note 69, at 21.

95. *See id.*

96. LORR, *supra* note 79, at 28.

97. The store was called the Piggly Wiggly because customers charged through like piggies to a trough. *See id.* at 29. This was just three years after the assembly line was first employed by Henry Ford. *Id.* at 29–30.

98. *Id.* at 30.

99. VICKI HOWARD, FROM MAIN STREET TO MALL: THE RISE AND FALL OF THE AMERICAN DEPARTMENT STORE 34 (2015).

100. King Kullen was the first supermarket, opening in 1930, and its "stores were warehouse-size shops." KIM HUMPHERY, SHELF LIFE: SUPERMARKETS AND THE CHANGING CULTURES OF CONSUMPTION 68 (1998).

101. *See* LORR, *supra* note 79, at 33.

grocery store in the United States was a supermarket.¹⁰² Over the course of a century, local shops gave way to department stores and supermarkets, then to discount chains and big-box retailers, then to suburban shopping malls. With these marketplace changes, the consumer experience was transformed.

Trademark law was developed to respond to this changed marketplace and consumer experience.¹⁰³ The first draft of what would eventually become the Lanham Act emerged in 1921,¹⁰⁴ just after branded products had started to become prevalent and firms were investing in brand names.¹⁰⁵ By the time the Lanham Act passed in 1946,¹⁰⁶ supermarkets had made their debut and the postwar households had an abundance of product choices. When search costs theory emerged in 1987, the American shopping mall and big box stores were opening everywhere.¹⁰⁷ Trademark law and theory was responsive to the changing marketplace from 1921 to 1987.

This brief history of shopping reveals that retail underwent a paradigm shift in the early twentieth century when consumer choice emerged. We are now experiencing the second paradigm shift in retail, in which consumer choice is once again becoming mediated—this time by AI. In the first shift, consumers shifted their trust in the shopkeeper to trademarks. In the second shift, consumers are shifting their reliance on trademarks to AI. The first shift produced trademark law. Trademark law should be responsive to the second shift.

B. *The Depreciation of Trademark's Information Function*

Search costs theory presupposes the stability of the information-carrying function of trademarks and assumes, without question, that consumers can trust in the trademark. Sometimes, however, consumers' trust in trademarks is misplaced.

102. *See id.*

103. *See* Mark Bartholomew, *Advertising and the Transformation of Trademark Law*, 38 N.M. L. REV. 1, 13–14 (2008) (explaining how the rise of advertising and mass production of goods changed the landscape of trademark law); Bone, *Hunting Goodwill*, *supra* note 29, at 576–78; Walter J. Derenberg, *Trade-Marks Ante Portas*, 52 YALE L.J. 829, 829 (1943) (“Today, with the tremendous development in modern methods of selling and advertising, trade-mark protection has become a matter of national concern . . .”).

104. Christine Haight Farley, *The Lost Unfair Competition Law*, 110 TRADEMARK REP. 739, 775 (2020).

105. Bartholomew, *supra* note 103, at 14 (“By 1920, however, advertising had changed. . . . Salesmanship of a brand became important; simple announcements of a product’s availability and content were insufficient.”).

106. Farley, *supra* note 104, at 775; *see* 15 U.S.C. § 1051.

107. MARLENE TARG BRILL, *AMERICA IN THE 1980S*, at 74 (2010).

Even before the age of algorithms, trademarks' "information transmission" theory¹⁰⁸ has been vulnerable to facts on the ground. As the marketplace has changed over time, the simple account search costs theory presents often fails to correspond with how trademarks are used. In numerous circumstances, trademarks fail to inform consumers about the source of the product, and in some, they actually conceal information that may be relevant to the consumer's purchasing decision. Nevertheless, the theory of trademarks prevails. It has artfully been stretched to fit new realities, but the consumers' interest is disserved.

Trademark doctrine reinforces the idea that when a consumer looks at a product and asks, "Who are you?," trademarks answer the question with, "It is me!"¹⁰⁹ The implication is that the consumer knows who that is.¹¹⁰ But exactly who does the consumer think it is?

Although trademark law holds that trademarks indicate to consumers the source of origin of products, the concept of "origin" has not been stable over time. One might expect origin to refer to the place from which the product came, and it may have, back when products were made in one place. Today, origin refers instead to the entity that owns the trademark.

As "origin" has been a slippery concept in trademark law, the Supreme Court found it necessary to address its meaning in 2003 in *Dastar Corp. v. Twentieth Century Fox Film Corp.*¹¹¹ The Court concluded that "the most natural understanding of the 'origin' of 'goods'—the source of wares—is the producer of the tangible product sold in the marketplace."¹¹² The producer of the tangible product, then, is to whom the consumer thinks the trademark refers.

Even the term "producer," however, has morphed over time in trademark law. In the beginning, the producer was the manufacturer of the tangible product. The trademark owner has not been the actual manufacturer of most consumer products, however, for several decades

108. See Bone, *Hunting Goodwill*, *supra* note 29, at 549 (coining this term).

109. In determining whether a term is generic, courts may employ the "who are you?/what are you?" test. See, e.g., *Off. Airline Guides, Inc. v. Goss*, 6 F.3d 1385, 1391 (9th Cir. 1993) (citation omitted) (noting that if the name answers the question "what are you?," it is deemed generic and unprotectable as a trademark, whereas if the name answers the question "who are you?," it may be protected as a trademark).

110. Consumers need not know who exactly is the source of the marked products so long as they understand the products emanate from a single, albeit anonymous, source. MCCARTHY, *supra* note 22, § 3:12.

111. 539 U.S. 23 (2003).

112. *Id.* at 31; see also GILSON LALONDE, *supra* note 30, § 1.03(1) ("The source or origin of goods under trademark law is 'the producer of the tangible product sold in the marketplace.'" (quoting *Dastar Corp.*, 539 U.S. at 31)).

now.¹¹³ The Ninth Circuit in *Siegel v. Chicken Delight, Inc.*¹¹⁴ explicitly recognized that “[t]he historical conception of a trade-mark as a strict emblem of source of the product to which it attaches has largely been abandoned.”¹¹⁵ Even in *Dastar*, the Court permitted that “[t]he concept might be *stretched* . . . to include not only the actual producer, but also the trademark owner who commissioned or assumed responsibility for (‘stood behind’) production of the physical product.”¹¹⁶

Search costs theorists have stretched the concept of origin even further, completely severing the relationship with any entity. Origin is instead a conceptual reference. For instance, in a companion article to Landes and Posner’s, Economides asserted that “[t]he consumer of NABISCO WHEAT THINS knows and cares little about source (manufacturer). Rather the consumer identifies the trademark with the features of the commodity.”¹¹⁷ Likewise, the *Siegel* court proclaimed “a new rationale for trade-marks” in which trademarks, rather than representing the producer, instead are “representations of product quality.”¹¹⁸ Thus, “source” now equates not to who makes it or where it comes from, but to the attributes of the product, such as quality. While consumers might associate a producer with certain attributes, the law has as its target something much more amorphous by shifting from manufacturer to product attributes.

The theory that trademarks inform consumers about quality has an important caveat, which is that the “quality” need not be high quality, only

113. For example, Levi’s has not manufactured its own products since 2003. *Levi Strauss Closes Last Two U.S. Plants*, NBC NEWS (Jan. 8, 2004, 1:42 PM PST), <https://www.nbcnews.com/id/wbna3909407> [<https://perma.cc/2HRE-K5T4>]. Likewise, an article on mattress firms quoted an industry representative as saying, “None of these guys, with a few rare exceptions, make their own mattresses.” Jasmine Wu, *There Are Now 175 Online Mattress Companies—and You Can’t Tell Them Apart*, CNBC (Aug. 19, 2019, 11:18 EDT), <https://www.cnbc.com/2019/08/18/there-are-now-175-online-mattress-companiesand-you-cant-tell-them-apart.html> [<https://perma.cc/9YF5-XCA8>].

114. 448 F.2d 43 (9th Cir. 1971).

115. *Id.* at 48.

116. *Dastar*, 539 U.S. at 31–32 (emphasis added). The Court, however, did put some limit on how far the concept of origin can be removed from the production of the goods:

the phrase ‘origin of goods’ is in our view incapable of connoting the person or entity that originated the ideas or communications that ‘goods’ embody or contain. Such an extension would not only stretch the text, but it would be out of accord with the history and purpose of the Lanham Act and inconsistent with precedent.

Id. at 32.

117. Economides, *supra* note 42, at 527. See generally *id.* (using all caps to indicate a trademark).

118. *Siegel*, 448 F.2d at 48.

consistent quality.¹¹⁹ Even this standard, however, imagines business relationships that do not equate with modern business practices. Where a trademark owner consistently contracts with the same manufacturer to produce the same product, ensuring consistent quality is feasible. Trademark owners today, however, contract with several different manufacturers. That same trademarked product may be produced by many additional manufacturers over time and even concurrently. Those manufacturers, in turn, may simultaneously produce that same product for competitor firms to be sold under their trademarks.¹²⁰ Even these manufacturers may subcontract in the manufacturing process. Although it may have initiated production and authorized the use of its trademark, the trademark owner has become so far removed from the actual production that quality control is possible only at the margins. Moreover, the source of origin of the ingredients or material in a product is increasingly difficult to ascertain in today's complex global supply chain and is more likely to be inconsistent.

Thus, in many cases, products are produced by other firms that are related to the trademark owner only by contract. Trademarks do not disclose this information to consumers. In fact, the trademark is more likely to obscure this information. Consumers may be dismayed to see the level of "control" that is hiding behind the trademark.

C. *Trademarks that Misinform Consumers*

Numerous realities of the marketplace have diminished the information function of trademarks. One reality that the information transmission theory ignores is the instability of the ownership of trademarks due to the volume of mergers and acquisitions (M&As), which has dramatically increased since the information transmission function was put forward.¹²¹ In some cases, the change in ownership can dramatically alter the connection with the trademarked product, such as where the owner is a

119. 3 MCCARTHY, *supra* note 22, § 18:55 ("It is important to keep in mind that 'quality control' does not necessarily mean that the licensed goods or services must be of 'high' quality, but merely of equal quality, whether that quality is high, low or middle.").

120. The mattress industry is illustrative. As one industry expert explained, "[mattress firms are] literally calling around to producers saying, 'we need a finished product and here's what we think it should look like.' Sometimes, they don't even know what they want it to look like." Wu, *supra* note 113.

121. See, e.g., Charles R. McManis, *Intellectual Property and International Mergers and Acquisitions*, 66 U. CIN. L. REV. 1283, 1283 (1998) ("[A] new wave of mergers is recasting the face of business in the United States."); MICHAEL E.S. FRANKEL & LARRY H. FORMAN, *MERGERS AND ACQUISITIONS BASICS: THE KEY STEPS OF ACQUISITIONS, DIVESTITURES, AND INVESTMENTS* 2 (2d ed. 2017) ("There is a long-term upward trend in both the volume and average deal size of acquisitions in the United States.").

huge conglomerate without a specific reputation for the trademarked product or where a firm engaged in unrelated business activities or has a wholly different reputation in that market.¹²² These M&As further attenuate the informational value of the trademark for the consumer.

The widespread assignment and licensing of trademarks also interferes with the transmission of useful information to consumers.¹²³ Trademark licensing was prohibited under common law and the Trademark Act of 1905 precisely because it caused the trademark to fail to accurately indicate the origin of the products.¹²⁴ The Lanham Act permits trademark licensing, but requires that trademarks be used by “related companies.”¹²⁵ Today, however, the trademark owner and the manufacturer are typically “related” only by the licensing agreement.

For trademarks to serve as an information source that improves market efficiency, they must enable consumers to overcome the asymmetry of information about the quality and characteristics of the product. Economides explains:

In many markets, sellers have much better information as to the unobservable features of a commodity for sale than the buyers. . . . Unobservable features, valued by the consumer, may be crucial determinants of the total value of the good. . . . The economic role of the trademark is to help the consumer identify the unobservable features of the trademarked product. This information is not provided to the consumer in an analytic form, such as an indication of size or a listing of ingredients, but rather

122. Some examples of this disconnection are Unilever’s purchase of BEN & JERRY’s, Clorox’s purchase of BURT’S BEES, and Colgate-Palmolive’s purchase of TOM’S OF MAINE. Kim Bhasin & Patricia Laya, *13 Small Eco Brands that Are Actually Owned by Giant Corporations*, BUS. INSIDER (Oct. 1, 2011, 12:34 PM PDT), <https://www.businessinsider.com/13-ethical-mom-and-pop-brands-that-are-actually-owned-by-giant-corporations-2011-10> [<https://perma.cc/8D2R-YY8S>]. Arguably, if the product remains the same, the trademark continues to serve its function, although who stands behind it has changed. While consumers can stop buying products that change in quality as a result of the new owners, the theory that the market will correct itself rests on a different conception of the consumer as hyper vigilant and introduces inefficiency.

123. See 2 JEROME GILSON, TRADEMARK PROTECTION AND PRACTICE § 6.01(2) (1998); Irene Calboli, *The Sunset of “Quality Control” in Modern Trademark Licensing*, 57 AM. U. L. REV. 341 (2007) (arguing that in today’s economy, quality control by trademark owners is increasing unsustainable).

124. See 3 MCCARTHY, *supra* note 22, § 18:39 (“In the early development of trademark law, trademarks were thought to identify only the physical source of the product or service in connection with which the mark was used.”); *Dawn Donut Co. v. Hart’s Food Stores, Inc.*, 267 F.2d 358, 366 (2d Cir. 1959) (“Prior to the passage of the Lanham Act many courts took the position that the licensing of a trademark separately from the business in connection with which it had been used worked an abandonment.”).

125. 15 U.S.C. § 1055 (“Where a registered mark or a mark sought to be registered is or may be used legitimately by related companies, such use shall inure to the benefit of the registrant or applicant for registration, and such use shall not affect the validity of such mark or of its registration . . .”).

in summary form, through a symbol which the consumer identifies with a specific combination of features.¹²⁶

Thus, the economic theory of trademarks is contingent on the fact of consistency in the manufacture of products, and also an argument that the legal protection of trademarks will ensure such consistency. The fact that today a trademark is affixed to several and varied products produced by multiple third parties is a repudiation of this theory.

The advent of wide brand extensions has produced increased trademark licensing. Brand extensions involve exploiting valuable brands to sell an increasingly unrelated array of products. According to the influential trademark treatise author J. Thomas McCarthy, the quality theory of source was directly responsible for these brand extension practices.¹²⁷ For example, Harley Davidson attaches its famous motorcycle brand to products as disparate as clothes, cake decorations, and toilet paper.¹²⁸ The manufacturers of these products likely designed these products and may even have initiated the licensing deal. Who then is the more accurate source of the products?

Retailers' private labels tell consumers that the retail chain is the source of origin of the products. At Trader Joe's almost every product is branded with the private label even though they come from various and fluctuating sources. Trader Joe's takes efforts to conceal this source information from the public.¹²⁹ Whether it is Kirkland or Amazon Basics, the products sold under these brands are simply renamed, pre-existing products. The private label therefore obscures source information.

The mattress industry provides an illustrative example.¹³⁰ Here, although there are many companies offering competitive products, and

126. Economides, *supra* note 42, at 526–27.

127. See 3 MCCARTHY, *supra* note 22, § 18:40 (“The quality theory led to the modern phenomenon of a trademark owner in market A licensing its mark to sometimes far-flung markets X, Y and Z.”).

128. Victoria Slind-Flor, *Money and Mayhem*, INTELL. ASSET MGMT. 15, 15 (2007) (mentioning Harley Davidson's brand extensions to clothing, beer, perfume, condoms, cribbage boards, furniture for children, coffee, beef jerky, bedding, Barbie dolls, wedding-cake toppers, and cake-decorating kits).

129. Vince Dixon, *What Brands Are Actually Behind Trader Joe's Snacks?*, EATER (Aug. 9, 2017, 1:16 PM EDT), <https://www.eater.com/2017/8/9/16099028/trader-joes-products> [<https://perma.cc/E5NT-WP3Y>] (“Trader Joe's and its suppliers all but swear to keep the agreement secret.”).

130. The eyewear industry is another example. One company, Luxottica, produces eighty percent of all branded eyeglasses and sunglasses. NAWAR AL-SAAD, TO UNDERSTAND THE WORLD UNDERSTAND ECONOMICS 60 (2022) (“[Luxottica] has bought almost all the major eyewear brands. However, they are still named differently. It creates an illusion in the consumer's mind that they have a variety of sunglasses to choose from, although they are all manufactured by one company.”). Luxottica owns or licenses to Oakley, Ray-Ban, and Versace, to name a few; it also runs LensCrafters, Pearle Vision, Sears Optical, Sunglass Hut, Target Optical, and EyeMed Vision Care. David Lazarus,

each firm itself offers numerous trademarked models for consumers to choose from, the trademarks on mattresses offer consumers little useful information and are often a barrier to accessing pertinent information. Although the profusion of trademarks insinuates a broad choice of producers, forty percent of all mattresses are produced by just one manufacturer.¹³¹ These manufactures supply the identical mattress to multiple retailers under different brand names exclusive to each retailer.¹³² Under this practice, Sealy Posturepedic Reyna Ridge mattresses may only be purchased at Macy's, while Bloomingdales and other retailers may sell identical mattresses, but under different trademarks. The deliberate consequence of this business practice is to prevent the consumer from comparison-shopping.¹³³ In this industry, trademarks perform the most minimal function of distinguishing one mattress from another within one retail outlet (although not between different retailers), but they fail to convey any information to consumers about the source of origin or characteristics of the products.

Consumers want to know about the firms with which they do business.¹³⁴ Today, consumers have expressed intense interest in knowing the level of social responsibility of the firms that offer them products. Evaluation of a firm's environmental, social, and corporate governance (ESG) has emerged as a means of expressing the public's concern over the general lack of transparency and accountability in today's corporations.¹³⁵ While trademarks could theoretically play an effective

How Badly Are We Being Ripped Off on Eyewear? Former Industry Execs Tell All, L.A. TIMES (Mar. 5, 2019, 5:00 AM PT), <https://www.latimes.com/business/lazarus/la-fi-lazarus-glasses-lenscrafters-luxottica-monopoly-20190305-story.html> [<https://perma.cc/3SUG-29FW>].

131. Wu, *supra* note 113.

132. See Steven Kurutz, *An Easy Choice? Dream On*, N.Y. TIMES (Oct. 8, 2014), <https://www.nytimes.com/2014/10/09/garden/how-to-find-the-best-mattress-in-the-maze-of-choices.html> (last visited Oct. 16, 2023) ("I'd go into store B and say, 'Do you have the Serta blah, blah, blah?' And the salesperson would say: 'I don't know. We may. But ours have different names.'" (quoting a consumer)); Timothy B. Lee, *Mattress Stores Want to Rip You Off: Here's How to Fight Back*, VOX (Aug. 10, 2015, 3:00 PM EDT), <https://www.vox.com/2015/3/5/8151607/mattress-buying-tips-savings> [<https://perma.cc/FQA9-5QEU>] ("[S]ome mattress manufacturers will give the same mattress different names in different stores.").

133. See Lee, *supra* note 132.

134. Cf. Laura R. Bradford, *Trademark Law and Agency Costs*, 55 INTELL. PROP. L. REV. 193, 203 (2015) (noting the "undersupply of clear information about sponsorship and affiliation" in the consumer market).

135. For instance, New York first introduced proposed legislation in 2021, later reintroduced in 2023, that would create environmental and social governance reporting requirements for multinational fashion retail sellers and manufacturers. See Fashion Sustainability and Social Accountability Act, A. 8352, 2021–2022 Assemb. (N.Y. 2021); Fashion Sustainability and Social Accountability Act, A. 4333, 2023–2024 Assemb. (N.Y. 2023). Among other disclosures, firms would be required to

role here, such as alerting consumers to stay away from certain branded products based on low ESG scores, it is instructive to consider how trademarks may have contributed to the lack of transparency in the first place. Today, many trademark owners have relinquished, at least to some extent, their capacity to control quality, which they had when they manufactured and sourced the raw materials themselves.

Not only do many trademark owners typically have no production capacity, but they also may not even employ designers or inventors. Some trademark owners are holding companies. When trademark owners neither produce, design, nor even commission the products sold under that trademark, it is difficult to determine what useful information the trademark is conveying to consumers. For example, a consumer might believe that the menswear company Robert Graham was the firm of a designer named Robert Graham. Instead, Robert Graham products are designed by unrelated firms that approach the company to suggest certain products.¹³⁶ If agreed to, those firms then source the product from a manufacturer, who will in turn source the materials from other manufacturers.¹³⁷ Robert Graham will sign several such contracts with various third parties.¹³⁸ Recalling the *Wizard of Oz*, trademarks can function like a curtain that shrouds a lone contract attorney.

There are numerous examples of trademarks that unsurprisingly, perhaps deliberately, lead consumers astray. For example, a consumer who cares about the environment and consuming healthy food may be attracted to Boca Burgers veggie burgers, imagining that the firm that stands behind that trademark is guided by those concerns. It is not. Boca Burgers is owned by the large tobacco firm Philip Morris.¹³⁹ Fashion

disclose supply chain mapping and social and environmental sustainability reports for “a minimum of” fifty percent of suppliers by volume. *Id.* These firms need a law to force them to learn just fifty percent of the supply chain for the goods that they affix with their trademark!

136. This business model was described to the author by the Robert Graham employee—who is not a lawyer—in charge of these decisions and contracts. Interview with Unnamed Robert Graham Employee, Robert Graham, in Washington, D.C. (June 11, 2014) (on file with author).

137. *Id.*

138. *Id.*

139. Dow Jones, *Kraft to Buy Boca Burger*, N.Y. TIMES, Jan. 19, 2000, at C13. Kraft Foods is a unit of Philip Morris. *Id.* Similarly, if a consumer has a bad rental car experience with National and therefore chooses Alamo the next time they rent a car, they may be disappointed to know that both brands are owned by Enterprise. See Andrew C. Taylor, *Enterprise’s Leader on How Integrating an Acquisition Transformed His Business*, HARV. BUS. REV. (Sept. 2013), <https://hbr.org/2013/09/enterprises-leader-on-how-integrating-an-acquisition-transformed-his-business> [<https://perma.cc/6TFR-FAL6>]. Another example of trademark misinformation was the news reports of people around the world expressively pouring out STOLICHNAYA vodka to protest Russia’s invasion of Ukraine. Claire Thornton, *Ditching Vodka to Protest Russia? You Might Be Surprised Where That Bottle Is Actually from*, USA TODAY (Mar. 4, 2022, 12:57 PM ET),

brands such as Christian Lacroix and Kate Spade likely indicate to consumers that the garments that carry the designer's names have been designed by or at least were designed under the authorization of the eponymous designer. In fact, in these cases and many others, the designer whose name the trademark bears is legally separate from the company.¹⁴⁰ Or consider the many examples of faux foreign trademarks like Ginsu knives (from Ohio, not Japan)¹⁴¹ or Häagen-Dazs, whose original ice cream cartons depicted a map of Denmark, when in fact the ice cream brand was founded in the Bronx as a made-up, European-sounding name.¹⁴² There is good reason to think that the misinformation trademarks convey is material. When consumers learned that Häagen-Dazs was not Danish, but came from the Bronx, sixty-eight percent decided not to buy it.¹⁴³ Even if we include the transmission of product characteristics under an information transmission theory approach, this function also does not work because consumers think they are getting product characteristics that they are not. In these scenarios, trademarks convey fraudulent assurance of product quality.

Even though the Supreme Court recently stated that “[t]he Lanham Act makes th[e] fact [of whether a mark’s use is serving a source-designation function] crucial[] in its effort to ensure that consumers can tell where

<https://www.usatoday.com/story/money/2022/02/26/vodka-russia-most-vodka-brands-u-s-arent-russian/6952976001/> [<https://perma.cc/FKJ5-YD6R>]. STOLICHNAYA, however, is made in Latvia—a NATO member country—and is owned by a firm in Luxembourg. *Id.* Ownership of the STOLICHNAYA trademark is the subject of a decades-long legal dispute. Macpherson Kelley, *Long-Running Dispute over Stolichnaya and Moskovskaya Vodka Marks Reignited*, LEXOLOGY (July 21, 2021), <https://www.lexology.com/library/detail.aspx?g=20e326cb-0da6-424b-bb0d-88fb09738931> [<https://perma.cc/SGR3-SP3E>]. As a result, the “origin” of STOLICHNAYA depends on the bottle.

140. See George C. Sciarino & Matthew D. Asbell, *The Designer Formerly Known as . . . : Intellectual Property Issues Arising from Personal Names as Fashion Brands*, 107 TRADEMARK REP. 1150 (2017) (citing the examples of designer labels KATE SPADE, CHRISTIAN LACROIX, DAVID J. PLINER, PAUL FRANK, JOSEPH ABOUD, CATHERINE MALANDRINO, and KAREN MILLEN); Yvette Joy Liebesman, *When Selling Your Personal Name Mark Extends to Selling Your Soul*, 83 TEMP. L. REV. 1, 27–31 (2010) (citing the examples of JOSEPH ABOUD and SIGRID OLSEN); Allison B. Pitzer, Comment, *Unfashionably Late: Protecting a Designer’s Identity After a Personal Name Becomes a Valuable Trademark*, 35 S. ILL. U. L.J. 309, 309–10, 318–24, 331–32 (2011) (citing examples of HERVE LEGER, Paolo Gucci of GUCCI, and JOSEPH ABOUD).

141. Andrew Guy Jr., *The Ginsu Guys Carved a Very Profitable Niche*, CHRON. (Mar. 13, 2005), <https://www.chron.com/life/article/The-Ginsu-guys-carved-a-very-profitable-niche-1934720.php> [<https://perma.cc/3R98-UXD8>]; see also *id.* (“‘Who would buy a set of knives made in Ohio? . . . We had to add some mystery.’” (quoting Barry Becher, co-founder of Ginsu)).

142. See Alison Spiegel, *Häagen-Dazs Doesn’t Come from Where You Think It Comes from*, HUFFPOST (Feb. 11, 2021), https://www.huffpost.com/entry/haagen-dazs-comes-from_n_7266208 [<https://perma.cc/FV9R-25MW>]. General Mills now owns Häagen-Dazs. *Id.*

143. Dan Nosowitz, *Häagen-Dazs Ice Cream Is from the Bronx—so What’s with the Name?*, ATLAS OBSCURA (Sept. 5, 2017), <https://www.atlasobscura.com/articles/haagen-dazs-fake-foreign-branding> [<https://perma.cc/6MH2-ACH2>].

goods come from,”¹⁴⁴ in so many cases, trademarks fail to inform consumers of the true provenance of the products they adorn. In fact, trademarks may serve to mask just how attenuated the origin of the product is from the trademark owner. In these instances, trademarks function less as information packets¹⁴⁵ than they do misinformation packets.

Economides, one of the economists that articulated the search cost theory to an enthusiastic trademark bar,¹⁴⁶ explained the benefit of trademarks from a consumer perspective. In Economides’ view, trademarks aid the consumer in a situation of information asymmetry.¹⁴⁷ The consumer can only see competitive products that appear similar; they have no way of ascertaining the products’ unobservable differences, which only the producers know.¹⁴⁸ By communicating these unobservable differences to consumers, trademarks, according to Economides, alleviate the information asymmetry.¹⁴⁹ However, the preceding examples demonstrate that information asymmetry still exists and has even been increased by the expansion of trademark protection. Search costs theory proponents, like Economides, fail to account for the information asymmetry that trademarks can engender. Unobservable differences between products also can be misrepresented to the consumer by means of the trademark.

III. TRADEMARKS IN THE AGE OF ALGORITHMS

A. *The Tech Challenge to Trademark Theory*

Thus far, neither scholarly critique nor the numerous inconsistent realities of the marketplace have been able to topple the ascendancy of search costs theory. However, with the AI-mediated marketplace, search costs theory has finally met its justifying limit. This paradigm shift in shopping fundamentally undercuts the rationale for trademarks.

If we protect trademarks because they are efficient at conveying accurate information, it means that the cost of protection outweighs the

144. *Jack Daniel’s Props., Inc. v. VIP Prods. LLC*, 599 U.S. ___, 143 S. Ct. 1578, 1593 (2023).

145. See Carter, *supra* note 29, at 759.

146. See Swann, *supra* note 60, at 1132 n.3 (indicating that the Economides article was funded by a grant from the United States Trademark Association (USTA), the precursor to INTA; that Landes and Posner’s article was derived from a lecture by Judge Posner at a USTA funded event; and that Economides, along with Landes and Posner, published articles in *The Trademark Reporter* in 1988).

147. See Economides, *supra* note 42, at 526.

148. See *id.*

149. See *id.*

alternative, which would be less efficient or less accurate. Grounded in the archetype of a consumer overwhelmed by a store shelf full of choices, such a tradeoff seems necessary. However, consumers today possess alternative means to efficiently research which product choice is optimal. They are no longer forced to rely solely on the informational shortcut provided by trademarks. For example, origin information can be far more accurately and efficiently disseminated through blockchain than trademarks.¹⁵⁰ Walmart is already using blockchain to tell consumers the provenance of its products.¹⁵¹ Other technological means may be more reliable than trademarks to trace products through the supply chain.¹⁵² An obvious one is that almost all consumers carry with them tiny computers that can access and process nearly all of the world's information in milliseconds. Landes and Posner could never have anticipated this technological advance. Consumers are now able to bring real research to bear on purchasing decisions.¹⁵³ Anything that can be known about the products we are considering travels with us on our mobile devices. Which dish soap works best on grease, is most environmentally friendly, smells best, is most economical, is made with fair labor standards, is made in America, etc.? A few clicks will provide all the answers.

The consumer search costs theory of trademarks is premised on consumers mainly relying on their personal experience with products—good and bad—as the primary source of information that drives future purchasing choices, with second-hand information and advertising pitches supplementing that first-hand experience.¹⁵⁴ Search costs theory is premised on the consumer approaching the store shelf with their choice already made and efficiently eyeing the trademark that corresponds to that choice. This strategy is efficient—so long as you can remember which soap your friend mentioned—but may not yield the optimal choice because of the paucity of inputs. Relying on such limited input, however, is now wholly obsolete. Today, massive amounts of customer reviews are

150. See, e.g., Ron Miller, *Walmart Is Betting on the Blockchain to Improve Food Safety*, TECHCRUNCH (Sept. 24, 2018, 12:00 PM PDT) <https://techcrunch.com/2018/09/24/walmart-is-betting-on-the-blockchain-to-improve-food-safety/> [<https://perma.cc/V9QN-ZL5H>] (“By placing a supply chain on the blockchain, it makes the process more traceable, transparent and fully digital.”).

151. See *id.*

152. See Margaret Chon, *Tracermarks: A Proposed Information Intervention*, 53 HOUS. L. REV. 421 (2015) (explaining how technology can provide traceable and verifiable source information to consumers).

153. Fifty-three percent of respondents answered the question, “How do you search for specific information on a product that you want to buy?,” with “Search engines (e.g., Google).” *Sources of Information About Products in the U.S. as of June 2023*, STATISTA (Aug. 25, 2023), <https://www.statista.com/forecasts/997051/sources-of-information-about-products-in-the-us#statisticContainer> (last visited Sept. 23, 2023).

154. Economides, *supra* note 42, at 532–33.

ubiquitous in the online environment. Why rely on the producer's recommendation or your own limited experience when there is easy access to thousands of customer reviews, third party ratings, and knowledgeable descriptions of attributes? Economist George Stigler, who is credited with advancing the concept of consumer search costs in economics, argued in 1961 that search costs are reduced when consumers pool information.¹⁵⁵ Stigler could not have foreseen how large that pool of information would become.

Now, the consumer who stands before a shelf of product choices may be armed with more relevant and accurate information to make that choice. In those moments, however, consumers will still need to rely on trademarks as the devices that enable them to match products with their researched selections. That is, for the internet-informed consumer, the trademark still serves as a distinguishing device. The information function of trademarks, however, is dramatically scaled back. The trademark merely identifies itself. If internet research had persuaded the consumer to buy the cheapest or most expensive choice, the price tag would have functioned in exactly the same way.

B. *The Disappearance of Trademarks from Search*

The marketplace has dramatically changed since the 1980s when the search costs theory was articulated.¹⁵⁶ The consumer experience is radically changing, with the pandemic accelerating these changes.¹⁵⁷ The change in where and how consumers shop¹⁵⁸ significantly bears on the function of trademarks in search. Consumers are shopping in the physical world at declining rates.¹⁵⁹ More consumers are doing more shopping

155. Cf. Stigler, *supra* note 60, at 216.

156. See, e.g., GILSON LALONDE, *supra* note 30, § 1.03(12) (illustrating the function of trademarks by describing a consumer selecting a product from a store shelf). How trademarks are advertised and function in the marketplace continues to be anachronistically described in trademark law even today. *Id.*

157. See Rae Yule Kim, *The Impact of COVID-19 on Consumers: Preparing for Digital Sales*, 48 IEEE ENG'G MGMT. REV. 212, 214 (2020).

158. See James Yang, *NPR/Marist Poll: Amazon Is a Colossus in a Nation of Shoppers*, NPR, <http://www.npr.org/about-npr/617470695/npr-marist-poll-amazon-is-a-colossus-in-a-nation-of-shoppers> [<https://perma.cc/V8P8-JLC9>] ("The habits of American shoppers are changing profoundly as more of them log on to buy things online. With one click or a voice command, things arrive magically at their doorsteps.").

159. See Simon Torkington, *The Pandemic Has Changed Consumer Behaviour Forever - and Online Shopping Looks Set to Stay*, WORLD ECON. F. (July 7, 2021), <http://www.weforum.org/agenda/2021/07/global-consumer-behaviour-trends-online-shopping/> [<https://perma.cc/XN3T-UHRA>]; Tamara Charm, Becca Coggins, Kelsey Robinson & Jamie Wilkie, *The Great Consumer Shift: Ten Charts That Show How US Shopping Behavior Is Changing*,

online¹⁶⁰ and this trend is predicted to continue.¹⁶¹ Millennials, who were projected to comprise one-third of total retail spending in 2021,¹⁶² and will eventually “have a greater impact on the economy than the Baby Boomer[s],”¹⁶³ are comfortable doing all of their shopping online. Likewise, Gen Zers are digital natives never having known a world of commerce that did not include the internet, social media, and mobile.¹⁶⁴

In the digital environment, making the optimal product choice is no longer an exercise in hunting for trademarks. According to the search costs model, consumers scan store shelves hunting for the package marked Crest. When shopping online, the consumer does not have to see Crest, or even remember Crest, but can follow clicks within their research to purchase the optimal toothpaste.

In the digital marketplace, trademarks are vanishing from search. In virtual retail, consumers cannot physically examine the product or its trademark and packaging. They see the physical product for the first time only after it has been purchased and delivered. Consequently, they may

MCKINSEY & CO (Aug. 4, 2020), <http://www.mckinsey.com/business-functions/marketing-and-sales/our-insights/the-great-consumer-shift-ten-charts-that-show-how-us-shopping-behavior-is-changing> [https://perma.cc/4L3G-YCR4]. In 2021, 78.5 percent of the U.S. population shopped online. *Penetration Rate of E-Commerce in the U.S. 2018–2028*, STATISTA, <http://www.statista.com/statistics/273958/digital-buyer-penetration-in-the-united-states/> (last visited Oct. 31, 2023).

160. E-commerce has been continuously increasing each year. See Seray Keskin, *19 New Ecommerce Statistics You Need to Know in 2023*, DRIP (May 24, 2022), <http://sleeknote.com/blog/e-commerce-statistics> [https://perma.cc/L44L-VSH2]; Alicia Phaneuf, *Ecommerce Statistics: Industry Benchmarks & Growth*, INSIDER INTEL (Jan. 8, 2022), <http://www.insiderintelligence.com/insights/ecommerce-industry-statistics/> [https://perma.cc/MY8H-TCZT] (“We expect US ecommerce sales will cross \$1 trillion for the first time in 2022.”). According to a Pew Research study in 2016, seventy-nine percent of U.S. consumers shop online. Aaron Smith & Monica Anderson, *Online Shopping and E-Commerce*, PEW RSCH. CTR. (Dec. 19, 2016), <http://www.pewresearch.org/internet/2016/12/19/online-shopping-and-e-commerce/> [https://perma.cc/4KLF-C4ZQ]. Over half shopped from their mobile phone, and fifteen percent purchased after clicking through on a link shared on social media. *Id.*

161. Recently, sales in e-commerce have increased five times faster compared to in-store retail. Kim, *supra* note 157, at 215. There is no indication that e-commerce is going to stop growing. See Kai Stübane, *Online Shopping Trends 2022: Pandemic-Era Habits Settle In*, THE FUTURE OF COM. (2022), <http://www.the-future-of-commerce.com/2022/02/04/online-shopping-trends-2022-pandemic-era-habits-settle-in/> [https://perma.cc/NU8X-AQ4B] (“Researchers at Statista estimate global e-commerce sales will top \$7 trillion dollars by 2025.”).

162. Iva Marinova, *18 Amazing Online Shopping Statistics 2023*, REV. 42, <https://review42.com/resources/online-shopping-statistics/> [https://perma.cc/6E6F-NXUP] (last updated May 20, 2023).

163. Susan M. Puwalski, *Millennials and E-Commerce: The Online Shopping and Purchasing Behavior of Millennials Attending College (2011)* (Ph.D. dissertation, Capella University) (ProQuest). Millennials are one-third larger than the Baby Boom generation. *Id.*

164. See Andris A. Zoltners, Sally E. Lorimer & Prabhakant Sinha, *How Sales Teams Can Thrive in a Digital World*, HARV. BUS. REV. (Feb. 18, 2020), <http://hbr.org/2020/02/how-sales-teams-can-thrive-in-a-digital-world> [https://perma.cc/3E3M-LQN9].

not have noticed how a trademark is displayed on the product or packaging when making the purchasing decision. Nontraditional trademarks such as smell, sound, taste, and tactile marks will therefore be of no assistance to consumers in making a purchase decision online. The trademark may be displayed virtually, but it will likely have a diminished impact. It may not be seen “affixed” to the product,¹⁶⁵ and the trademark may be displayed only in plain typeface, in black, and in small font, as is other information such as price.

More significantly, trademarks may not even appear in search. Source information may instead be contained in source code, invisible to consumers. When consumers find products on their phones and on social media, they rely on this digital information.¹⁶⁶ Before, trademarks were a stand-in for product research. Now, purchases actually result from research. A consumer may arrive at a trademarked product, but the path to that product was not through the assistance of the trademark. According to the search costs theory, trademarks “promote and secure business reputation and goodwill by securing a mnemonic device between products and source.”¹⁶⁷ Numerous technologies, however, relieve consumers of the need of mnemonics to make or repeat a satisfactory purchase. A simple one is the “order it again” feature that many e-retailers offer. In such cases, consumers may not even notice what brand of candles or charging cables they have been buying.

In the digital marketplace, many of today’s sellers have already recognized the irrelevance of trademarks. Online shopping practices reveal that in many instances brand matters less than other factors and may not even matter at all. Delivery charges and delivery dates have become a major factor in purchasing decisions.¹⁶⁸ For instance, for

165. Affixing the trademark to the good is required for trademark protection. 15 U.S.C. § 1127(1)(A) (“For purposes of this chapter, a mark shall be deemed to be in use in commerce . . . on goods when . . . it is placed in any manner on the goods or their containers or the displays associated therewith or on the tags or labels affixed thereto, or if the nature of the goods makes such placement impracticable, then on documents associated with the goods or their sale . . .”). Use in commerce is required for trademark protection. *See id.*

166. Sixty-one percent of retail traffic comes from mobile devices. Keskin, *supra* note 160.

167. Dan L. Burk & Brett H. McDonnell, *Trademarks and the Boundaries of the Firm*, 51 WM. & MARY L. REV. 345, 351 (2009).

168. Lauren Freedman *The Shopper Speaks: Shipping and Delivery Satisfaction Numbers Stand Strong*, DIGIT. COM. 360 (Aug. 30, 2023), <https://www.digitalcommerce360.com/2023/08/30/shipping-delivery-satisfaction> [https://perma.cc/5JPZ-BL7K]; *see* Martin Joerss, Florian Neuhaus & Jürgen Schröder, *How Customer Demands Are Reshaping Last-Mile Delivery*, MCKINSEY & CO. (Oct. 19, 2016) <https://www.mckinsey.com/industries/travel-logistics-and-infrastructure/our-insights/how-customer-demands-are-reshaping-last-mile-delivery> [https://perma.cc/5VRZ-NHGM] (indicating that delivery options, speed, and perceived quality are key differentiators amongst large e-commerce

consumers shopping on Amazon, searchable product characteristics and verified customer satisfaction may drive their search.¹⁶⁹ Placement in search results, however, may be more important than any other factor.¹⁷⁰

Consumers on Amazon may find in their top search results products sold under unrecognizable marks like BSTOEM and ZGGCD.¹⁷¹ Such seemingly random strings of letters are an increasing presence on Amazon.¹⁷² These “pseudo-brands” are neither selected nor deployed to function as trademarks in the traditional sense. Instead, they are adopted solely to optimize placement in search results.¹⁷³ For example, the FRETREE trademark appears in the listings of a large range of unrelated products such as “‘ice cream scoops,’ ‘animal-activated pet feeders’ and ‘camping grills.’”¹⁷⁴ The trademark owner also owns the registered trademarks DRALEGEND and CORLITEC, which appear in random product listings including “alarm clocks, flashlights, blowtorches and yoga mats.”¹⁷⁵

Unpronounceable and meaningless in every language, these marks have been acquired only to respond to the idiosyncrasies of the Amazon platform. They are not used in marketing and usually do not even appear on the products. They are a creature of Amazon policies and algorithms.¹⁷⁶

players); see also Laura Lane, *The Future of Social Commerce and Marketplaces: eConsultancy Report*, CHANNELADVISOR (Sept. 21, 2020), <http://www.channeladvisor.com/uk/blog/marketplaces/the-future-of-social-commerce-and-marketplaces-econsultancy-report/> [https://perma.cc/SWM7-3TDB] (“Eighty-three percent of consumers claim convenience is more important than it was five years ago.”).

169. See, e.g., Jie Yang, Rongting Zhang & Vanessa Murdock, *Exploring the Reasons Behind Customer Purchase Decisions*, AMAZON SCI. (Mar. 10, 2020), <http://www.amazon.science/blog/exploring-the-reasons-behind-customer-purchase-decisions> [https://perma.cc/4YAA-Z6R6] (reporting on a study finding that “perceived quality is highly correlated with price”); Rosie Murphy, *Local Consumer Review Survey 2020*, BRIGHTLOCAL (Dec. 9, 2020), <http://www.brightlocal.com/research/local-consumer-review-survey-2020/> [https://perma.cc/MN2Y-M52Y] (finding that seventy-nine percent of respondents said they trusted online reviews as much as the recommendations of family and friends).

170. A whopping “49% of shoppers aged 27-40 say they always or frequently/often buy the first product listed on an Amazon” search engine results page. FEEDVISOR, *THE 2021 AMAZON CONSUMER BEHAVIOR REPORT* 4, http://s3.amazonaws.com/media.mediapost.com/uploads/2021_Consumer_Behavior_Report.pdf [https://perma.cc/WE8E-NT7R].

171. John Herrman, *All Your Favorite Brands, from BSTOEM to ZGGCD*, N.Y. TIMES (Feb. 11, 2020), <http://nyti.ms/38uhKvi> (last visited Nov. 3, 2023).

172. *Id.*

173. See *id.*

174. *Id.*

175. *Id.*

176. Note, *Fanciful Failures: Keeping Nonsense Marks off the Trademark Register*, 134 HARV. L. REV. 1804, 1805 (2021). Amazon’s policies have caused a surge in Chinese trademark filings in the

In an effort to be responsive to brand owners' concerns over counterfeit goods on the platform, Amazon developed the "Brand Registry" program, which requires a U.S. trademark registration.¹⁷⁷ Participants in the program receive the benefit of appearing higher up in search results and earlier access to consumer reviews, both of which significantly impact sales.¹⁷⁸ The fastest and cheapest way to acquire a trademark registration is to apply for a made-up word that has never been used by anyone before and has no meaning that may create an obstacle to registration.¹⁷⁹ A trademark is only a ticket to entry into the Amazon marketplace; it has no other value to these sellers or to consumers.

The fact that products sold under marks like ZGGCD can be "Best Sellers" on Amazon demonstrates the comparatively low value of brand to other criteria such as search result, price, shipping, and product reviews.¹⁸⁰ According to Professors Barton Beebe and Jeanne Fromer's study of a random sample of applications for marks for apparel filed at the United States Patent and Trademark Office (USPTO) in 2017, 44.4% (162 of 365) applications from China—which account for about a third of all applications—involved a "nonsense word that is unpronounceable in English and that the applicant indicated has no meaning in any other language."¹⁸¹

The rationale for trademarks is undermined by the marked change in the way consumers now make purchasing decisions. Internet research capabilities, of course, are just one small step away from the consumer purchasing decision of a bygone era. Other developments, however, expose the consumer search costs model as simply inapt and anachronistic. Rather than a range of choices presented to a consumer who

United States Patent and Trademark Office. *See id.* at 1804. Consequently and perversely, U.S. public servants spend inordinate time assisting foreign sellers and Amazon in reaching financial success. *Id.* at 1814.

177. Herrman, *supra* note 171.

178. *Id.*

179. *Id.* Amazon also offers its "IP Accelerator Program, through which sellers can join Amazon's Brand Registry [even] before obtaining a trademark" by applying for trademark registration with the help of an approved law firm. *Id.*

180. *See, e.g.*, Cleo Levin, *Made-for-Amazon Brand Names Are Getting Ridiculously Surreal*, SLATE (Oct. 14, 2020, 10:20 AM), <http://slate.com/technology/2020/10/amazon-brand-names-pukemark-demonlick-china.html> [<https://perma.cc/5LXP-HU4L>] ("There is also a startling amount of competition among inexpensive, third-party clothing brands. If you search 'women's long sleeve shirts' on Amazon, you get more than 80,000 results, the overwhelming majority of which do not come from recognizable brands. Even if Pukemark does not have its own website, or even a dedicated landing page on Amazon, the name can still function as something to catch people's attention long enough to get them to click through and buy an \$8 shirt.")

181. Barton Beebe & Jeanne C. Fromer, *Fake Trademark Specimens: An Empirical Analysis*, 120 COLUM. L. REV. F. 217, 232 (2020).

must then select among them, today numerous apps and services are at consumers' disposal to direct them to the product of their desires. If a consumer sees something they like the looks of, either on screen or in the wild, they can, with the aid of their phones and QR codes, barcodes, SKU numbers, model numbers, etc., discover where to purchase such a product. In such a transaction, the trademark does not play an identification role in either the search for, or the selection of, the product. Increasingly, consumers can rely on image alone to lead them to the product they want. For instance, if a consumer sees a chair they like in a photo on Pinterest, Pinterest Lens will connect them with chairs for sale that are visually similar.¹⁸² These technologies use machine learning to provide recommendations for similar items online. In all of these searches, trademarks do not aid the consumer in finding the product, although they may help consumers choose if there are multiple accurate results.

Another technology that substitutes for trademarks in the marketplace is near-field communication (NFC). Chips in phones exchange data when they come into contact with objects with NFC tags, which can be added to almost any tangible product.¹⁸³ NFC tags embed product information

182. See *Pinterest Lens*, PINTEREST, <http://help.pinterest.com/en/article/pinterest-lens> [<https://perma.cc/5US4-2LM3>]. For a human-built analog, try Worn on TV, which provides information about where to buy clothing and accessories that looks like what celebrities have worn. See *WORN ON TV*, <http://wornontv.net/> [<https://perma.cc/S8WS-GPCX>]. For example, Target partners with Pinterest to use photos that customers have pinned to suggest visually similar items that it sells. Tricia McKinnon & Ben Rudolph, *20 Innovative Examples of Artificial Intelligence in Retail*, INDIGO9 DIGIT. (May 19, 2020), <http://www.indigo9digital.com/blog/artificialintelligence> [<https://perma.cc/S6WP-LCXB>]. Google Lens, Bing's Visual Search, and Amazon's StyleSnap all enable consumers to use image search to find products. See *GOOGLE LENS*, <http://lens.google/> [<https://perma.cc/T7J9-43NZ>]; Raghav Haran, *Bing Now Lets You Search for an Object Within an Image*, SINGLE GRAIN, <http://www.singlegrain.com/news/bing-now-lets-search-object-within-image/> [<https://perma.cc/3FG3-3HAP>]; AMAZON, <http://www.amazon.com/stylesnap> [<https://perma.cc/VEN2-655E>]. CamFind offers a similar app-based tool. See *Visual Search*, CAMFIND, <http://camfindapp.com/> [<https://perma.cc/YK9P-WAKR>].

183. Peter Dahlström & David Edelman, *The Coming Era of 'On-Demand' Marketing*, MCKINSEY & CO. (Apr. 1, 2013), <http://www.mckinsey.com/business-functions/marketing-and-sales/our-insights/the-coming-era-of-on-demand-marketing> [<https://perma.cc/J29L-LSCY>]. Google offers "window shopping," which allows users to search for products using general terms. See *Window Shopping*, CHROME WEB STORE, <http://chrome.google.com/webstore/detail/window-shopping/oifaapodofedmholecgkpiplenfeiehdo?hl=en> [<https://perma.cc/UNC5-FTEX>]. For example, "snow boots" produces a visual feed of products that can be further filtered by prompted terms such as "women," "fur," "lace up," and "free returns." *Id.* (follow "Window Shopping" hyperlink; then download the "Window Shopping" app; then search "snow boots"). Results include customer reviews, rankings, and price. *Id.* Brands can be search terms, but a highly specific search can be accomplished without them. *Id.* Unlike shopping at an online retailer, the results are pulled from twenty-four billion product listings across the internet. Antonio G. Di Benedetto, *Google Expands Shopping Searches with Lens and In-Store Inventory Checks*, VERGE (Sept. 29, 2021, 10:30 AM PDT), <https://www.theverge.com/2021/9/29/22696646/google-shopping-lens-search-inventory-check-ios-chrome> [<https://perma.cc/T9BY-CZB2>].

in goods, enabling consumers to research a product when they encounter it in the world. Rather than conducting remote research, consumers explore products through an interactive shopping experience. For example, noticing a headset, a consumer can tap it with their phone for information.¹⁸⁴ The phone might then suggest taking a selfie to see what the consumer looks like wearing various colors or models of the headset. The consumer can then post these images on social media for peer feedback. The consumer may then receive a customized offer such as bundling a service with the purchase of the headset, relying on current subscription data. As the consumer wears this new headset, it may detect when the phones of those who reacted to the consumer's post are near and those individuals may then be offered a special deal. At no time in this scenario does a trademark appear in the strategic marketing or in the consumer's search. While a trademark may still serve as a referent, these are not trademark searches; these are product searches, and they are remarkably efficient.

C. *The Advent of Adtech*

The story in the introduction may sound futuristic and dystopian, but data collectors and advertisers are using all of those technologies now. The deployment of AI in the digital marketplace was rapidly accelerated by the pandemic as more consumers turned to online purchases.¹⁸⁵ Big data, AI, machine learning, and deep learning have already transformed the marketplace, and even more changes are on the way.

The advertising technology industry known as "adtech" is a vast and growing industry.¹⁸⁶ AI is predicted to have a \$40 billion effect on marketing by 2025.¹⁸⁷ The fact that the internet has largely been built on an advertising-based business model has given rise to surveillance

184. Dahlström & Edelman, *supra* note 183.

185. Joe McKendrick, *AI Adoption Skyrocketed Over the Last 18 Months*, HARV. BUS. REV. (Sept. 27, 2021), <https://hbr.org/2021/09/ai-adoption-skyrocketed-over-the-last-18-months> [<https://perma.cc/W249-NYU4>]; see also Kim, *supra* note 157, at 213; Stübane, *supra* note 161.

186. The digital advertising market was expected to grow 30.5 percent in 2021 to revenues of \$491 billion globally. Alistair Gray, *Three Tech Giants Control Half of Advertising Outside China*, FIN. TIMES (Dec. 7, 2021), <http://www.ft.com/content/bcbe8674-060f-4298-aab8-91e40e00c3f2> (last visited Oct. 13, 2023). Google, Facebook, and Amazon have doubled their share of ad revenues in the past five years. *Id.* Amazon's advertising business alone generated \$31 billion in revenue in 2021. Farhad Manjoo, *The Rise of Big Tech May Just Be Starting*, N.Y. TIMES (Feb. 16, 2022), <http://www.nytimes.com/2022/02/16/opinion/big-tech-stock-market.html> (last visited Nov. 3, 2023).

187. Vance Reavie, *Do You Know the Difference Between Data Analytics and AI Machine Learning?*, FORBES (Aug. 1, 2018, 7:00 AM EDT), <http://www.forbes.com/sites/forbesagencycouncil/2018/08/01/do-you-know-the-difference-between-data-analytics-and-ai-machine-learning/#5c50edac5878> [<https://perma.cc/EV88-NEKF>].

capitalism, where personal data is the commodity exchanged.¹⁸⁸ Recent antitrust cases brought against big tech firms have revealed how online advertising is the key component.¹⁸⁹ For example, Alphabet, Google's parent company with a revenue of \$257 billion in 2021,¹⁹⁰ has a market capitalization that exceeds the GDP of all but eleven countries.¹⁹¹ Significantly, more than eighty percent of that revenue comes from targeted advertising. Based on their revenue models, Google and Facebook might be thought of as advertising companies.¹⁹² When tech companies this large become advertising companies, advertising changes.

Remember the story about how Target marketed baby products to a family who did not yet know—as Target did—that their teenage daughter was pregnant?¹⁹³ That was more than ten years ago! With little notice,

188. See generally SHOSHANA ZUBOFF, *THE AGE OF SURVEILLANCE CAPITALISM* 8–10 (2019); TIM WU, *THE ATTENTION MERCHANTS: THE EPIC SCRAMBLE TO GET INSIDE OUR HEADS* (2016). Ironically, in an academic paper Google founders Larry Page and Sergey Brin famously warned that “advertising funded search engines will be inherently biased towards the advertisers and away from the needs of the consumers.” Sergey Brin & Lawrence Page, *Reprint of: The Anatomy of a Large-Scale Hypertextual Web Search Engine*, 56 *COMPUT. NETWORKS* 3825, app. A, at 3832 (2012).

189. See Complaint at 1, *Texas v. Google LLC*, No. 4:20-cv-957 (E.D. Tex. Dec. 16, 2020). Google boasts that “more daily transactions are made on [Google’s] AdX than on the NYSE and the NASDAQ combined.” Second Amended Complaint at 2, *In re Google Digit. Advert. Antitrust Litig.*, No. 1:21-md-03010-PKC (S.D.N.Y. Oct. 22, 2021). Nearly thirty percent of Google’s gross digital advertising revenue flows from its demand side platforms. Complaint at 90–91, *Dep’t of Just. v. Google LLC*, No. 1:23-cv-00108 (E.D. Va. Jan. 24, 2023).

190. Press Release, Alphabet Inc., *Alphabet Announces Fourth Quarter and Fiscal Year 2021 Results* (Feb. 1, 2022), https://abc.xyz/investor/static/pdf/2021Q4_alphabet_earnings_release.pdf [<https://perma.cc/58FR-MGZE>].

191. To illustrate, the market valuation of Alphabet Inc. is \$1.97 trillion. Akash Sriram & Subrat Patnaik, *Alphabet Eyes \$2 Trillion Value After Blowout Results*, *REUTERS* (Feb. 3, 2022, 5:47 AM PST), <http://www.reuters.com/business/alphabet-inches-closer-2-trln-market-value-after-blowout-results-2022-02-02/> [<https://perma.cc/PUB5-RGQ9>]. If it were a country, it would be the twelfth richest country, just below Brazil and Canada. See Omri Wallach, *The World’s Tech Giants, Compared to the Size of Economies*, *VISUAL CAPITALIST* (July 7, 2021), <https://www.visualcapitalist.com/the-tech-giants-worth-compared-economies-countries/> [<https://perma.cc/6KA9-85V8>].

192. See Complaint at 2–3, *Texas v. Google LLC*, No. 4:20-cv-957 (E.D. Tex. Dec. 16, 2020) (“[Google’s] entire business model is targeted advertising . . .” (emphasis in original)); Hannah Murphy, *Facebook Patents Reveal How It Intends to Cash In on Metaverse*, *FIN. TIMES* (Jan. 18, 2022), <http://www.ft.com/content/76d40aac-034e-4e0b-95eb-c5d34146f647> (referencing Meta’s “existing \$85bn-a-year ad-based business model”) (last visited Nov. 7, 2023). The recent dispute between Apple and Facebook over Apple’s new privacy settings, which disrupts Facebook’s data-harvesting ad services, demonstrates how dependent these firms have become on their ad revenue models. See Patience Haggin, Keach Hagey & Sam Schechner, *Apple’s Privacy Change Will Hit Facebook’s Core Ad Business. Here’s How.*, *WALL ST. J.* (Jan. 29, 2021, 11:49 AM ET), <http://www.wsj.com/articles/apples-privacy-change-will-hit-facebooks-core-ad-business-heres-how-11611938750> (last visited Nov. 7, 2023).

193. Charles Duhigg, *How Companies Learn Your Secrets*, *N.Y. TIMES* (Feb. 16, 2012), <http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html> (last visited Oct. 14, 2023).

Target began to employ predictive analytics more than twenty years ago.¹⁹⁴ The volume and sophistication of the patents that have been filed in this space indicate how much more advanced data-driven advertising has developed since then.¹⁹⁵ For instance, Meta's patents reveal how it can present hyper-targeted advertising in augmented reality.¹⁹⁶

Marketers today know vast amounts about us and have increasingly sophisticated means of using this knowledge to predict our purchasing choices. Online monitoring, the Internet of Things, wearable technologies, and smart homes have already crept into our lives. These technologies enable a cycle of data collection, machine learning, and targeted solicitation.¹⁹⁷ The technologies enable firms to know what we need, want, and like better than we do.¹⁹⁸ As a result, consumers no longer have to waste time with shopping lists or hunting around for the product that fulfills their material desires. The bots will do it for us; and they will do it

194. *Id.*

195. See Murphy, *supra* note 192 (noting Meta's recent USPTO applications "reveal that Meta has patented multiple technologies that wield users' biometric data in order to help power what the user sees and ensure their digital avatars are animated realistically" but "also indicate how the Silicon Valley group intends to cash in on its virtual world, with hyper-targeted advertising and sponsored content"). For example, Facebook/Meta was issued U.S. Patent No. 11,239,399 on February 1, 2022, for a method of personalized advertising and U.S. Patent No. 11,244,996 on February 8, 2022, for a method for tracking users' facial expressions. Google was issued U.S. Patent No. 11,250,679 on February 15, 2022, that will enable cookie-less ad targeting.

196. *See id.*

197. They also come at a significant cost to our privacy, as well as our economic and mental health. *See, e.g.,* ZUBOFF, *supra* note 188, at 8 (decrying how surveillance capitalism profits from the capture and analysis of behavioral data); Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995, 1028 (2014) ("The question of what constitutes a privacy violation is generally tied to control over personal information, with the logical consequence that increased collection and processing of data is usually linked to a greater privacy threat."); Kate Crawford & Jason Schultz, *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*, 55 B.C. L. REV. 93, 96 (2014) (explaining how consumers are harmed by collection of personal information).

198. *See* BEN PRING, EUAN DAVIS & VICTORIA BOLAND, COGNIZANT, 21 MARKETING JOBS OF THE FUTURE 20 (2019), <http://www.thecognizant.com/site/assets/files/2168/21-marketing-jobs-of-the-future-codex4428.pdf> [<https://perma.cc/ZNP6-4YKK>] ("[C]ustomers no longer have to think about what they'd like to buy or do because the bots do all the hard work."); *id.* at 44 ("In the connected world of smart things, it won't be people making recommendations to their friends, but a combination of people and algorithms driven by AI and machine learning."); Utpal M. Dholakia, *The Perils of Algorithm-Based Marketing*, HARV. BUS. REV. (June 17, 2015), <http://hbr.org/2015/06/the-perils-of-algorithm-based-marketing> [<https://perma.cc/U4HY-6R7D>] ("Algorithms help marketers utilize customer-specific knowledge — demographics, previous behavior, fellow customers' choices — to craft customized offers and deliver them, often in real time."); Jennifer Valentino-DeVries & Jeremy Singer-Vine, *They Know What You're Shopping for*, WALL ST. J. (Dec. 7, 2012), <http://www.wsj.com/articles/SB10001424127887324784404578143144132736214> (last visited Nov. 7, 2023) (discussing how comprehensively and precisely consumers are tracked online).

better and faster.¹⁹⁹ Gone are the days when consumers had to ask their friends for product recommendations. What their personal contacts are repeatedly buying, not returning, and positively reviewing may already be some of the inputs that go into the bot's product suggestion. A friend's recommendation of a laundry detergent is utterly subpar by comparison. Did the friend recall the correct trademark? How many different products did the friend test? Was the friend's evaluation biased by advertising or other positive associations with the brand? Are the friend's washing machine, water, and laundry habits sufficiently similar? Machines do not need to rely on limited or inaccurate information as consumers formerly did when they have vast stores of data on which to rely.²⁰⁰ Search costs theory, like most economic theories, is premised on rational actors—consumers who make rational choices. Responding to targeted ads that are based on reliable data may be more rational than relying on memory, a friend's recommendation, or a non-personalized advertisement.

Formerly, personal shopping assistants, concierge services, and stylists were a luxury few could afford. Now, these services can be available to everyone as digital personal assistants—powered by algorithms, AI, and machine learning—crawl the internet to find products to recommend.²⁰¹ Consumer data and algorithms enable marketers to successfully suggest products that targeted consumers may have either not known existed, or realized they wanted. The deep knowledge these marketers have means that consumers may be solicited without first indicating that they are in the market. Consumers may not be searching for products, but marketers are finding them anyway. Consumers will soon learn that the products they want or need will come to them without their needing to search.

199. See, e.g., McKinnon & Rudolph, *supra* note 182 (highlighting Walmart's Jetblack service, which enables customers to text their shopping needs to Walmart, which will populate a shopping cart with suggested items).

200. See, e.g., *id.* (noting that Kroger Grocery's AI-powered EDGE (Enhanced Display for Grocery Environment) app causes digital displays to change as customers approach to make personalized recommendations and offers).

201. *Id.* at 44. For instance, Amazon Assistant, a now-discontinued Chrome extension, enabled Amazon to make personalized product recommendations as consumers browsed other websites. See Rory Mellon, *Amazon Is Killing One of Its Most Underrated Shopping Features*, TOM'S GUIDE (Mar. 24, 2023), <https://www.tomsguide.com/news/amazon-is-killing-one-of-its-most-underrated-features-this-sucks> [https://perma.cc/GB8F-QSY8]; see also Daniel Allen, *What Is Amazon Assistant?*, ANDROID POLICE (Oct. 10, 2022), <https://www.androidpolice.com/what-is-amazon-assistant/> [https://perma.cc/87YU-5YNP] (“In addition to giving individual product price comparisons, Amazon Assistant also does its best to guess what you're shopping for on a competitor's site and show you what Amazon offers. It seems to accomplish this by spying on what you've put in the search bar.”).

Consumers are increasingly delegating the responsibility of making purchasing choices out of convenience.²⁰² Product recommendation engines have become ubiquitous on platforms such as Amazon, Google, Netflix, and Facebook. Even as early as 2013, thirty-five percent of consumer purchases on Amazon and seventy-five percent of Netflix streaming resulted from recommendations based on algorithms.²⁰³ Just as Netflix and Spotify use AI to learn what you like to make personalized recommendations, marketers use AI to recommend purchases. Consumers are recognizing that algorithms can make purchasing decisions faster and better than they can.

As retail has become increasingly virtual, marketing ever more data-driven, and purchasing decisions driven by AI, predictive analytics are changing consumers' shopping experience and their behavior. Consumers are becoming more comfortable shopping on their AI-powered home devices and voice-controlled devices. It is a short distance from a consumer saying, "Alexa, buy cinnamon," to Alexa announcing, "I have added an item to your shopping cart that will be perfect for your nephew's ninth birthday next week." Brand loyalty will be an important input, for a time.²⁰⁴ AI, however, can easily recognize how brand is merely a proxy for other characteristics that may prove to be more reliable inputs over time. In advertising, predictive analytics, or "predictive intelligence," relies on an ecosystem of data harvesting and AI to know what the consumer wants or needs before the consumer does.²⁰⁵

The consumer's ultimate satisfaction with a recommended product will reflect the accuracy of the program's algorithms, and will be fed back into the machine's learning. This satisfaction may or may not result in a consumer's brand loyalty, but it will likely result in loyalty to the personal assistant bot. In this way, producers may be increasingly more interested in conveying product information to machines than to consumers.

202. This delegation is, in some ways, the natural next step from relying on machines for repeat purchases. Examples include subscription orders from various firms managed online, the Amazon Dash button that allows consumers to re-order products at the press of a button, and Alexa voice-activated repeat purchases. *E.g.*, *Welcome to Amazon Customer Service*, AMAZON, <https://www.amazon.com/gp/help/customer/display.html> (last visited Nov. 19, 2023).

203. Ian MacKenzie, Chris Meyer & Steve Noble, *How Retailers Can Keep Up with Consumers*, MCKINSEY & CO. (Oct. 1, 2013), <http://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers> [<https://perma.cc/N9MN-5LRR>].

204. *See, e.g.*, *Shop with Alexa*, AMAZON, <https://www.amazon.com/alexa-shopping-hub> [<https://perma.cc/C2BF-PM5U>] (offering examples of how Alexa can help consumers repeat purchases, create grocery lists, and search for popular brands by using past purchasing behavior data).

205. *See, e.g.*, ACCENTURE, *THE POST-DIGITAL ERA IS UPON US: ARE YOU READY FOR WHAT'S NEXT?* 5 (2019), http://www.accenture.com/_acnmedia/pdf-94/accenture-techvision-2019-tech-trends-report.pdf [<https://perma.cc/D5LD-7AF9>] (describing Sam's Club's use of machine learning to mine purchase data and create a "smart" auto filled grocery shopping lists for customers).

Machine-to-machine communication offers accuracy and efficiency unparalleled by the search costs theory. In this AI marketplace, where machines communicate to machines, trademarks are becoming superfluous.

Search costs theory holds that trademarks enable consumer agency in purchasing decisions, but consumers are now willingly delegating certain aspects—and sometimes all—of this agency. Now AI drives purchasing decisions as it does the shopping for consumers. AI is assisting consumers' purchasing decisions, but it is also supplanting human choice. These AI-driven purchasing decisions are highly efficient and accurate. Search is costless and unlike consumers, AI is not prone to trademark deception. A consumer may confuse Coke and Koke, but AI can identify the difference. Concern for search costs is not relevant in the digital marketplace.

Trademark law has thus far failed to acknowledge the fundamental disruptions to the marketplace brought about by adtech. The marketplace that trademark law hinges on is becoming increasingly anachronistic. The current rationale of trademarks has no correspondence to a marketplace where consumers are micro-targeted with data-driven algorithms to suggest purchases that may be precisely what the consumer wants, even though they did not know they desired it.

Adtech combines targeted advertising—showing ads only to consumers with defined characteristics²⁰⁶—with tracking technology, which enables advertisers to target users as they move around the internet.²⁰⁷ Online, targeted advertising is facilitated by platforms, such as Google and Facebook, that aggregate data on their users and make it available to advertisers. Micro-targeted advertising uses the same technologies but is hyper-targeted based on increased stores of data and

206. See Randal C. Picker, *Online Advertising, Identity and Privacy* 17–18 (U. Chi. L. Sch., John M. Olin L. & Econ. Working Paper No. 475, 2009).

207. Adam Tanner, *How Ads Follow You from Phone to Desktop to Tablet*, MIT TECH. REV. (July 1, 2015), <http://www.technologyreview.com/2015/07/01/167251/how-ads-follow-you-from-phone-to-desktop-to-tablet/> [<https://perma.cc/28CY-M4G8>].

AI.²⁰⁸ Based on predictive analytics, micro-targeted ads lead to a higher purchase rate.²⁰⁹

Consider the ability of Meta, Facebook's parent company, to target consumers. Facebook tracks demographics, behavioral attributes, and interest categories for every user in the United States. It learns users' interests from clicks, likes, websites visited, and apps and services used.²¹⁰ It also gathers data based on its users' activity across Meta companies and products, including information scraped from Facebook and Instagram profiles and websites that users have logged into with Facebook.²¹¹ In addition, it gathers data from other websites and apps that use Facebook's "business tools."²¹² If a consumer looks at a shirt on a retailer's site, this information can be sent back to Facebook and used by any advertiser with which it contracts.²¹³ Facebook also gathers geo-location data on users that enables advertisers to target consumers based on their movement.²¹⁴ Additionally, Facebook invites advertisers to contribute to its store of data by uploading specific customer information such as phone numbers, email addresses, home addresses, and birthdays into its business tools. Facebook then associates this personal identifying information with a Facebook profile, enabling it to learn more about its users.²¹⁵ Facebook uses all this data to train its algorithms to suggest subsets of users to target with specific ads.

Facebook has created specialized targeted advertising products for marketers. For example, Facebook's look-alike audience targeting feature

208. Oana Barbu, *Advertising, Microtargeting and Social Media*, 163 *PROCEDIA SOC. & BEHAVIORAL SCI.* 44, 44 (2014) (defining micro-targeting as "a way to successfully create personalized messages or offers, correctly estimate of their impact (in regards to sub-grouping) and delivery directly to individuals" (citing TOM AGAN, *SILENT MARKETING: MICRO-TARGETING* (2007), <http://adage.com/images/random/microtarget031207.pdf> (last visited Nov. 18, 2023))); *see also* Ashwin Machanavajjhala, Aleksandra Korolova & Atish Das Sarma, *Personalized Social Recommendations - Accurate or Private?*, 4 *VLDB ENDOWMENT* 440 (2011); Dokyun Lee, Kartik Hosanagar & Harikesh S. Nair, *Advertising Content and Consumer Engagement on Social Media: Evidence from Facebook* (Working Paper, 2015).

209. *See* MacKenzie et al., *supra* note 203 (listing Amazon, Google, Netflix, and Facebook among the list of companies seeing higher purchase rates); YAN LAU, *BUREAU OF ECON., FTC, ECONOMIC ISSUES: A BRIEF PRIMER ON THE ECONOMICS OF TARGETED ADVERTISING* 5 (2020).

210. Till Speicher, Muhammad Ali, Giridhari Venkatadri, Filipe Nunes Ribeiro, George Arvanitakis, Fabricio Benevenuto, Krishna P. Gummadi, Patrick Loiseau & Alan Mislove, *Potential for Discrimination in Online Targeted Advertising*, 81 *PROC. MACH. LEARNING RSCH.* 1, 3, 7 (2018).

211. *About Facebook Ads*, FACEBOOK, <http://www.facebook.com/ads/about> [<https://perma.cc/2CXV-QXVK>].

212. *Id.*

213. *Id.*

214. *Id.* Facebook also monitors where its users connect to the internet and where they use their phones. *Id.*

215. *Id.*; *see also* Speicher et al., *supra* note 210, at 4–5.

can build a population of target consumers that corresponds with a list of actual customers an advertiser provides.²¹⁶ In addition, Facebook's marketing application programming interface enables advertisers to input a piece of text and receive a suggested list of other correlating attributes.²¹⁷ These terms may then lead to other matched terms that can help an advertiser more precisely target an audience.

Firms work together to combine numerous data points on individual consumers. A post on Twitter reported in the press explains how this works.²¹⁸ By visiting someone's home, a consumer may then receive ads for the toothpaste brand purchased by the homeowner. The advertiser made the connection not because a smart speaker heard the two mention the brand or tracked the consumers' online search for the brand. Instead, this was possible because many apps collect data from our phones such as our unique device IDs and our locations. Data aggregators pay to pull data from everywhere they can. One source of rich data trails is consumers' purchases that are linked to loyalty programs, such as a grocery store or frequent flyer program. Aggregators buy these datasets and match them with datasets from various different firms by consumers' phone numbers and email addresses. The GPS location of our phones can also be monitored. These locations can be mapped and compared to other peoples' phone to reverse engineer our contacts. At this point, the interests and predilections of our contacts can be cross-referenced against our browsing and purchase histories. Upon all this data, products are suggested, such as the toothpaste used by one of our contacts.

The vast amounts of data these firms are armed with tells them which targets offer the best return on investment and enables them to micro-target these consumers.²¹⁹ Marketers know that sales are influenced by peer suggestions, and that some peers are more influential than others in

216. *How to Use Custom or Lookalike Audiences*, META, <http://www.facebook.com/business/help/572787736078838?id=176276233019487> [<https://perma.cc/79F7-MR33>].

217. Speicher et al., *supra* note 210, at 8–9. For instance, an advertiser can input “Fox” and receive the suggestion “The Sean Hannity Show,” which on Facebook will correspond with a subset that is a ninety-five percent match. *Id.* at 10–11.

218. Robert G. Reeve (@RobertGReeve), TWITTER (May 24, 2021, 8:32 PM), <http://twitter.com/robertgreeve/status/1397032784703655938> [<https://perma.cc/B9T4-3M6W>].

219. One of the important new positions in adtech is the data ethnographer, who combs data to reveal nuanced insights about consumers to employ in targeting advertising. See Siddharth Venkataramakrishnan, *What Marketers Need to Understand About Their Industry's New Technology*, FIN. TIMES (Feb. 25, 2020), <http://www.ft.com/content/834cc9ce-3ba2-11ea-b84f-a62c46f39bc2> [<https://perma.cc/B8KM-5LNE>]; see also Graeme Wood, *Anthropology Inc.*, ATLANTIC (Mar. 2013), <http://www.theatlantic.com/magazine/archive/2013/03/anthropology-inc/309218/> (last visited Nov. 7, 2023) (noting that Microsoft employs the second largest number of anthropologists after the United States government).

prompting sales.²²⁰ For instance, Instagram users can tag products in the photos they post that enable others to immediately purchase that product simply by clicking it.²²¹ What might seem like an exercise of choice—to rely on a select peer—may in fact have been invisibly orchestrated by algorithms. Significantly, such purchases are made without the search assistance of trademarks, which may not be mentioned or seen.

The correspondence between trademarks and search costs is diminishing for both consumers and marketers. Formerly, trademarks were useful proxies for marketers for particular market segments and demographics, such as income, location, or age. Data driven marketing, however, has exposed them as crude proxies for audience.²²² Narrowly targeted ads are clicked as much as 670% more than ordinary ones.²²³ Why waste ads on sweeping audiences—known as “spray and pray”—when a marketer can reach, on a mass scale, the precise consumers who will respond through micro-targeting.²²⁴ Today, ad buyers “can aim their ads at as few as 20 of the 1.5 billion daily users of [Facebook].”²²⁵

D. *The Disappearance of Search and Ads*

These technologies have completely transformed the very nature of how sellers sell. These are not your father’s banner ads. If an advertiser is guessing at what a consumer might click and purchase, it had better present this suggestion in the form of an advertisement so that the consumer understands why the information has been presented. However, when an advertiser knows with some degree of confidence that the

220. See generally Ravi Bapna & Akhmed Umyarov, *Do Your Online Friends Make You Pay? A Randomized Field Experiment on Peer Influence in Online Social Networks*, 61 MGMT. SCI. 1902 (2015) (studying empirically how peer influence impacts decision-making).

221. *Get the Latest from Instagram*, INSTAGRAM (Oct. 6, 2021), <https://business.instagram.com/blog/instagram-shopping-product-tags-customers> [<https://perma.cc/FD2J-3XQ8>] (“If you have a business or creator account, have uploaded products to your catalog and enabled Instagram Shopping, you can tag the products that are featured in your content so customers can shop in the moment of discovery.”).

222. One marketing firm drives this point home by listing consumer attributes that advertisers may want to target, such as “[h]as 2 dogs,” “[v]egetarian,” and “[l]arge collection of watches.” *Personalized Creative for Brands*, FUSION92, <https://media.fusion92.com/personalized-creative-for-brands/> [<https://perma.cc/JHD3-6Q8A>].

223. OSBORNE CLARKE, INTERNATIONAL ONLINE BEHAVIOURAL ADVERTISING SURVEY 2010, at 7 (2010).

224. See Natasha Singer, *Your Online Attention, Bought in an Instant*, N.Y. TIMES (Nov. 17, 2012), <https://www.nytimes.com/2012/11/18/technology/your-online-attention-bought-in-an-instant-by-advertisers.html> (last visited Nov. 5, 2023).

225. Natasha Singer, *‘Weaponized Ad Technology’: Facebook’s Moneymaker Gets a Critical Eye*, N.Y. TIMES (Aug. 16, 2018), <https://www.nytimes.com/2018/08/16/technology/facebook-microtargeting-advertising.html> (last visited Nov. 5, 2023).

consumer will buy the suggested product, the best strategy may be to sneak in the suggestion without announcing itself as an ad. For example, such ads appear on Facebook just as the user's friends' updates do.

Not only do product suggestions look different than ads, but algorithms can also make suggestions to consumers even when they are not shopping.²²⁶ The marketer need not wait for the consumer to approach the marketplace when it can unobtrusively reach consumers wherever they are and at any time. As result of the transformation of online selling, consumers are changing the way they shop. Consumers have fewer reasons to go to physical stores now that sophisticated digital shopping comes to them. They do not even need to go to a store's website; simply checking one's phone can initiate shopping. We are in the era of "ubiquitous shopping," where to be online using any device is to be shopping.²²⁷

In the "integrated digital marketplace"²²⁸ that is now embedded in our daily lives, shopping is not only pervasive, but also invisible. New surveillance technologies seamlessly integrate advertising into our daily lives. Consumers can shop through both video and voice. A 2018 retail industry rag gushed, "[v]oice recognition technologies are . . . taking the commercial battleground right into the heart of consumers' homes."²²⁹

Ads are not triggered by search, but are instead triggered by our mere presence online.²³⁰ Consumers are being monitored across a vast ecosystem of smart devices using, among other technologies, geo-location, voice recognition, facial recognition, image classification, video

226. In 2015, the Federal Trade Commission issued guidance on "native advertising" in an effort to protect consumers from ads that do not appear to be ads. See FED. TRADE COMM'N, ENFORCEMENT POLICY STATEMENT ON DECEPTIVELY FORMATTED ADVERTISEMENTS, https://www.ftc.gov/system/files/documents/public_statements/896923/151222deceptiveenforcemen nt.pdf [<https://perma.cc/4X6Q-MBBU>].

227. See RETAIL INDUS. LEADERS ASS'N (R)TECH CTR. FOR INNOVATION & ACCENTURE, *Delivering for the New Customer: The Move to Ubiquitous & Ultra-Personal Shopping* (2018), <https://rilastagamedia.blob.core.windows.net/rila-web/rila.web/media/media/pdfs/reports/accnture-report-the-changing-consumer-and-the-new-definition-of-retail.pdf?ext=.pdf> [<https://perma.cc/Z4T7-GB2N>].

228. *Id.*

229. Michelle Grujin, *Ubiquitous Shopping: Why This Trend Is Set to Transform the Next Phase of Retail – Inside Retail*, WINDOWSWEAR (Dec. 15, 2018), <https://www.windowswear.com/ubiquitous-shopping-why-this-trend-is-set-to-transform-the-next-phase-of-retail-inside-retail/> [<https://perma.cc/NP8X-9FEE>].

230. See Bingjie Liu & Lewen Wei, *Machine Gaze in Online Behavioral Targeting: The Effects of Algorithmic Human Likeness on Social Presence and Social Influence*, 124 COMPUTS. HUM. BEHAV., Nov. 2021, at 1, 3 (dubbing "machine gaze" as the ability for technology to personalize algorithms in online behavioral targeting from presence online).

analysis, and even biometric data from wearable technology.²³¹ Not only what we click, but even our natural, physical gestures can trigger ads while we are connected.²³² When an ad finds you, it may have relied on immense stores of information about you including your preferences, motivations, and relationships.

Advertisers are endeavoring to dissolve the line between commerce and content. “Frictionless transactions,” those seamless purchasing interactions via digital wallets and in-app and contact-less payments, lead to more instantaneous purchases. Beyond faster checkout, marketers remove friction by giving ads the illusion of passive discovery. Formerly, ads aimed to drive traffic from content sites to retail sites. Now, ads do not navigate away from content, but seek to enable shopping within content.²³³ Whereas traditional product placement in movies and television could only suggest *future* purchases, content itself has now become immediately shoppable.²³⁴ Armed with knowledge of who the viewer is, marketers can personalize the product depicted, showing them something they are bound to want and can buy instantly, or algorithms can suggest content that contains the products that are a match for them. Marketers aim to present “authentic” content based on knowledge of the consumer that “provides solutions for the consumer’s individual lifestyle

231. Yogesh K. Dwivedi, Laurie Hughes, Abdullah M. Baabdullah, Samuel Ribeiro-Navarrete, Mihalis Giannakis, Mutaz M. Al-Debei, Denis Dennehy, Bhimaraya Metri, Dimitrios Buhalis, Christy M.K. Cheung, Kieran Conboy, Ronan Doyle, Rameshwar Dubey, Vincent Dutot, Reto Felix, D.P. Goyal, Anders Gustafsson, Chris Hinsch, Ikram Jebabli, Marijn Janssen, Young-Gab Kim, Jooyoung Kim, Stefan Koos, David Kreps, Nir Kshetri, Vikram Kumar, Keng-Boon Ooi, Savvas Papagiannidis, Ilias O. Pappas, Ariana Polyviou, Sang-Min Park, Neeraj Pandey, Maciel M. Queiroz, Ramakrishnan Raman, Philipp A. Rauschnabel, Anuragini Shirish, Marianna Sigala, Konstantina Spanaki, Garry Wei-Han Tan, Manoj Kumar Tiwari, Giampaolo Viglia & Samuel Fosso Wamba, *Metaverse Beyond the Hype: Multidisciplinary Perspectives on Emerging Challenges, Opportunities, and Agenda for Research, Practice and Policy*, INT’L. J. INFO. MGMT., Oct. 2022, at 1.

232. “Ubiquitous computing,” or “ambient computing,” seeks to make human movement and gestures intelligible to computers through radar sensors. See Mark Wilson, *Google Is Designing Computers that Respect Your Personal Space*, FAST CO. (Mar. 1, 2022), <https://www.fastcompany.com/90725730/google-is-designing-computers-that-respect-your-personal-space> [<https://perma.cc/5BMX-SGQ7>]; see also Nicole Nguyen, *New Amazon Echo Show 10 Review: Alexa Has Got Its Eye on You*, WALL ST. J. (Feb. 24, 2021, 8:00 AM ET), <https://www.wsj.com/articles/amazon-echo-show-10-review-alexa-has-got-its-eye-on-you-11614171608> (last visited Nov. 5, 2023).

233. See, e.g., Umberto Torrielli, *The Post-Cookie Future: Driving Outcomes Through Context*, BLOOMBERG SPONSORED CONTENT, <https://sponsored.bloomberg.com/article/business-reporter/the-post-cookie-future-driving-outcomes-through-context> (last visited Nov. 18, 2023) (mentioning the resurgence of contextual advertising).

234. For example, AMC Networks has introduced its TEAL (Technology Enabled Audience Led) program that enables users to shop from television shows. Parker Herren, *AMC Networks Brings Shoppable Ads to the Upfronts*, AD AGE (Apr. 6, 2022), <https://adage.com/article/special-report-tv-upfront/amc-networks-brings-shoppable-ads-upfronts/2410681> (last visited Nov. 5, 2023).

and comfortably prompts action.”²³⁵ The goal is to present consumers with curated shoppable content wherever they happen to be rather than waiting until they have decided to shop.

Targeted ads now not only target the right consumer, but they also target the right time to present the product. Digital marketing now utilizes biometric data to monitor and analyze consumers’ emotional states.²³⁶ Emotional advertising²³⁷ is substantially more effective when it can target consumers at precisely the right moment, when they are in the emotional state to receive the suggestion.²³⁸ By tapping into a consumer’s emotional state in real time, marketers know just when consumers feel confident, motivated, vulnerable, or in search of comfort.

Consumers used to watch ads; now, ads watch them. Now that the marketer has near-perfect information about the consumer, the consumer has little need of the informational value of the trademark. In this world, trademarks play an increasingly small role. This is not the world to which search costs theory was addressed.

IV. REEVALUATING TRADEMARKS IN AN AI-DRIVEN MARKETPLACE

The AI-driven marketplace undercuts trademark law’s rationale. The search costs theory posits that trademarks are worthy of protection because of their value in efficiently assisting consumers to make informed purchasing decisions.²³⁹ Consequently, search costs theory is an argument that the social utility of efficiently conveying product information to consumers outweighs any anticompetitive effect of monopolizing information.²⁴⁰ Yet, as compared with trademarks, adtech provides consumers with higher levels of both accuracy and efficiency. By making this comparison, however, this Article does not argue that consumers are

235. Lane, *supra* note 168.

236. One up-and-coming job in marketing is the “Neuro A/B Tester” who “report[s] on the brain activity of customers when they come into physical contact with our brand during experiential and sensory campaigns.” PRINGET AL., *supra* note 198, at 12. Another is the “Mood & Empathy Manager” who “engineer[s] the right feelings at precisely the right point, so customers decide to engage longer.” *Id.* at 16. Additionally, the “Science Liaison/Bio-marketing Specialist” uses biometric data collected from consumer’s physiological responses to deliver “highly personalized brand communications.” *Id.* at 34.

237. See generally Peter Noel Murray, *How Emotions Influence What We Buy*, PSYCH. TODAY (Feb. 26, 2013), <https://www.psychologytoday.com/us/blog/inside-the-consumer-mind/201302/how-emotions-influence-what-we-buy> [<https://perma.cc/5SV7-6G7C>].

238. See Murphy, *supra* note 192 (“[E]ye gaze direction and pupil activity may implicitly contain information about a user’s interests and emotional state . . .”).

239. See *supra* section I.C.

240. See *supra* Part I.

better off with adtech than without it. Although we are rightly skeptical of adtech, trademarks are justified by the contention that they serve and protect consumers. Entrenched in this contention is the notion that consumers need trademarks to make informed choices.

The algorithmic world, however, is increasingly eliminating human agency in shopping. While consumers formerly made decisions about what they needed or wanted, when to shop, and which product to purchase, adtech has replaced these choices with a single decision: to buy now or not. The challenges that these technologies present to our settled understanding of the value of trademarks force us to rethink the cost-benefit analysis of protecting them. Perversely, adtech lays bare how trademarks can be used against consumers' interests. Though adtech is designed to operate stealthily, as a phenomenon its objective of manipulating consumer behavior is transparent. Juxtaposed with trademarks, as this Article has done, adtech also makes the disservice trademarks sometimes do to consumers more transparent than it has been in the past. Consumer choice has always been a site of manipulation.

In the context of adtech, it has become more evident that consumers are not adequately protected when product information is not accurate and they lack the ability to detect the inaccuracy.²⁴¹ At a minimum, if the information conveyance function can now be performed without the aid of trademarks, at least in some cases, the diminished value of trademarks should be accounted for in the law. To analyze the cost-benefit analysis of trademark protection, the benefit—the informational value of trademarks—must be reconsidered. And the value of the information conveyance performed by trademarks must be evaluated in light of the current availability of alternative means.²⁴² The search costs theory assumes the unavailability of alternative means to address the asymmetry of information between producers and consumers.²⁴³ Now, however,

241. Consumer protection may be achieved through false advertising law and administrative regulation by the Federal Trade Commission. Federal false advertising law protects consumers against advertising that “misrepresents the nature, characteristics, qualities, or geographic origin of . . . goods, services, or commercial activities.” 15 U.S.C. § 1125(a). Under the Federal Trade Commission Act, the FTC has authority to prohibit any practices in the marketplace deemed to be “unfair or deceptive.” 15 U.S.C. § 45. Michael Grynberg, *More Than IP: Trademark Among the Consumer Information Laws*, 55 WM. & MARY L. REV. 1429, 1433 (2014) (“[T]rademark law might benefit from ‘offloading’ some of its expanding scope to other consumer information regimes.”).

242. For example, Glynn Lunney called for a reevaluation of the cost-benefit analysis of trademark law generally by articulating three factors to be considered: (1) the extent to which the trademark conveys otherwise unavailable information; (2) the value of the information; its materiality; and (3) the availability of alternative means of conveying this information. Lunney, *supra* note 29, at 435.

243. See, e.g., S. REP. NO. 79-1333, at 3 (1946) (“[One] purpose underlying any trade-mark statute . . . is to protect the public so it may be confident that, in purchasing a product bearing a

consumers can utilize an arsenal of digital technologies to efficiently connect them with products that meet their desires. The need to rely on trademarks to convey information about products is therefore diminished. In some cases, the trademark may be helpful to consumers, while in others it may be superfluous or even an interference in making an optimal choice. In any event, these technological alternatives dramatically change the value of trademarks for consumers and so should change the cost-benefit analysis of protecting them. The search costs rationale for protecting trademarks is no longer descriptively accurate in many, if not most, cases.

Scholars have warned of the significant statutory and judicial expansion of trademark protections since the 1980s.²⁴⁴ The consensus view in the academy is that there is now too much protection.²⁴⁵ In the face of these mounting critiques, however, trademark rights have continued to be ratcheted up.²⁴⁶ The search costs theory has immunized trademark law from these critiques. The protection of trademarks is now out of step with the functional value of trademarks. Today, and more so in the future, trademarks may offer no informational value, no search efficiency, and may even obfuscate such informational values. The correspondence between trademarks and the search costs rationale is lost.

particular trade-mark which it favorably knows, it will get the product which it asks for and wants to get.”).

244. See, e.g., Lunney, *supra* note 29, at 371 (“This expansion has encompassed both the recognition of new trademark subject matter and a more complete bundle of ownership rights” (footnotes omitted)); Kenneth L. Port, *The Congressional Expansion of American Trademark Law: A Civil Law System in the Making*, 35 WAKE FOREST L. REV. 827, 830 (2000) [hereinafter Port, *Congressional Expansion*] (“None of the existing social, economic, or legal justifications supporting American trademark law encourage, let alone tolerate, such [trademark] expansion.”); Jessica Litman, *Breakfast with Batman: The Public Interest in the Advertising Age*, 108 YALE L.J. 1717, 1721–28 (1999) (criticizing the rationale of trademark law in various lawsuits alongside evolving trends in advertising and public interest); Lemley, *supra* note 57, at 1697–98 (arguing that courts replaced the traditional rationale for trademark law with new rationales favoring trademark owners); Glynn S. Lunney, Jr., *Trademark’s Judicial De-Evolution: Why Courts Get Trademark Cases Wrong Repeatedly*, 106 CALIF. L. REV. 1195, 1197 (2018) (warning that courts expanded trademark law far too broadly and caused anticompetitive outcomes); David J. Franklyn, *Owning Words in Cyberspace: The Accidental Trademark Regime*, 2001 WIS. L. REV. 1251, 1252 (2001) (“That courts have expanded exclusivity rights in favor of mark owners and gradually loosened restrictions on alienation is illustrative of the fact that trademark law has become increasingly propertized.”); Kenneth L. Port, *The Expansion Trajectory: Trademark Jurisprudence in the Modern Age*, 92 J. PAT. & TRADEMARK OFF. SOC’Y 474, 474 (2010) (“Congress has enlarged the trademark right at the behest of special interests without paying attention to the consequences.”).

245. See e.g., Port, *Congressional Expansion*, *supra* note 244, at 829 (“[Trademark law] development is not only inconsistent with the common law origins of the United States’ trademark system, but it also rests on questionable constitutional ground.”).

246. McGeeveran, *supra* note 51, at 51 (“Trademark rights have expanded dramatically over time from consumer-oriented safeguards against the diversion of customers to comprehensive protectors of brand identity.”).

To begin recalibrating trademark protections in light of the diminished functional value trademarks now offer vis-à-vis other available technologies, one must account for another change in the marketplace that has not yet disturbed the rationale of trademark protection. Although we may acknowledge it in our language, the shift in concept from trademark to brand has only had minimal impact on the law. The consequence of understanding trademarks as the sought-after good in and of themselves, will be particularly evident in the virtual world.

A. *The Morphed Function of Trademarks*

Although this Article predicts that in the future trademarks will be superfluous,²⁴⁷ there is one remaining function that trademarks may serve. Increasingly, the trademark, or “brand,” functions not as a packet of information that travels with the product, but as the product itself. Consumers desire the product because they want to own the brand. The brand is the commodity exchanged.

As the informational value of trademarks has diminished over time, for many trademarks, their commercial value has grown. The world’s most valuable brands are steadily increasing in value out of proportion to inflation.²⁴⁸ It is no coincidence that as trademarks’ commercial value has increased, so have their legal protections.²⁴⁹

247. See also Lemley, *supra* note 57, at 1688 (asserting that courts are protecting marks “as things valuable in and of themselves” and that recent developments “threaten to stretch the rationale of trademark law beyond all limits”); Litman, *supra* note 244, at 1718 (“Legal protection for trade symbols, in the absence of confusion, disserves competition and thus the consumer.”); Lunney, *supra* note 29, at 371 (stating that the “shift from viewing a trademark as a source of information about a product, to viewing the trademark as the product, has sharply changed the emphasis and context in which trademark’s traditional themes have played out”).

248. INTERBRAND, BEST GLOBAL BRANDS 2021, at 15 (“The combined value of the top 100 brands increased from \$2,326,491 million in 2020 to \$2,667,524 million in 2021, an overall increase of 15%.”).

249. See Lunney, *supra* note 29, at 371 (criticizing the “substantial and ongoing expansion of trademark protection”); Deborah R. Gerhardt, *Consumer Investment in Trademarks*, 88 N.C. L. REV. 427, 429 (2010) (“Trademark law has lost its way.”); Lemley, *supra* note 57, at 1688 (“Courts protect trademark owners against uses that would not have been infringements even a few years ago and protect as trademarks things that would not have received such protection in the past.”); Mark P. McKenna, *The Normative Foundations of Trademark Law*, 82 NOTRE DAME L. REV. 1839, 1840 (2007) [hereinafter McKenna, *Normative Foundations*] (“Trademark law unquestionably covers much more ground now than it did at the beginning of the twentieth century, and many of its more recent developments are fundamentally inconsistent with the normative goals of traditional trademark protection.”); Bone, *Trademark Infringement*, *supra* note 51, at 1330–36; Rebecca Tushnet, *Registering Disagreement: Registration in Modern American Trademark Law*, 130 HARV. L. REV. 868, 869 (2017) (“Courts too readily find too many acts to be infringing even when they’re harmless or actually useful to consumers.”).

The terminology has shifted from “trademark” to “brand” in recognition of their changed value in the marketplace. Historically, in the domain of law, lawyers, judges, and scholars spoke of “trademarks,” while trademark owners and marketing professionals referred to “brands.”²⁵⁰ Recently however, even in the law the discussion is increasingly about brands. Both the International Trademark Association (INTA)²⁵¹ and the World Intellectual Property Organization (WIPO)²⁵² have embraced this shift in terminology, despite the reference to trademarks and intellectual property in their names.

Though the legal community may have been slow to pivot from trademark to brand,²⁵³ the business world and the marketplace witnessed the phenomenon of the brand in the late eighties.²⁵⁴ Therefore, even as the Chicago School was describing the function of the trademark, trademark law was already in the midst of a major shift from the role it had played from the early twentieth century.

There has always been somewhat of a curious mismatch between marketing scholarship and practice, on the one hand, and trademark law, on the other.²⁵⁵ In marketing literature, a “brand” is defined as “a

250. Deven R. Desai, *From Trademarks to Brands*, 64 FLA. L. REV. 981, 985 (2012) (“While trademark law sees trademarks and brand as synonymous, brand scholarship and practice recognize that they are not.”).

251. For instance, in 2011, INTA described itself as “a not-for-profit membership association dedicated to the support and advancement of *trademarks* and related intellectual property.” *About INTA: Overview*, INT’L TRADEMARK ASS’N (emphasis added), <https://web.archive.org/web/20110110013639/http://www.inta.org/About/Pages/Overview.aspx> (last visited Nov. 18, 2023). In contrast, it currently states, “[w]e’re a global network of *brand* owners and professionals dedicated to supporting trademarks and related intellectual property. Explore how we support *brands* through our work in advocacy, events, and resources” INT’L TRADEMARK ASS’N (emphasis added), <https://www.inta.org/> [<https://perma.cc/QG2T-6YVV>].

252. In its 2013 report, the WIPO stated that “[f]rom its humble beginning as an identifier of origin, branding has evolved into a sophisticated business tool employing professionals as diverse as data analysts, lawyers, linguists, graphic artists, psychologists and celebrity actors.” WORLD INTEL. PROP. ORG., 2013 WORLD INTELLECTUAL PROPERTY REPORT: BRANDS – REPUTATION AND IMAGE IN THE GLOBAL MARKETPLACE 3 (2013), https://www.wipo.int/edocs/pubdocs/en/intproperty/944/wipo_pub_944_2013.pdf [<https://perma.cc/858Q-5W7L>] [hereinafter 2013 WORLD INTELLECTUAL PROPERTY REPORT].

253. Because the theoretical foundation of trademark law has always bewildered courts, it is certainly possible to find old trademark cases that seem to have recognized the capacity of brands. For instance, a 1901 British decision described a trademark as “the attractive force which brings in custom.” *Inland Revenue Comm’rs v. Muller & Co’s Margarine, Ltd.*, [1900–03] All ER Rep. 413 (HL).

254. See, e.g., DAVID A. AAKER, *MANAGING BRAND EQUITY: CAPITALIZING ON THE VALUE OF A BRAND NAME* 5 (1991) (describing Proctor & Gamble’s enormous success in branding and advertising).

255. See Jeremy N. Sheff, *The Ethics of Unbranding*, 21 FORDHAM INTEL. PROP. MEDIA & ENT. L.J. 983, 989–90 (2011) (noting ethical approaches and distinctions in trademark law doctrine); Sheff,

‘reputational asset’ which has been ‘developed over time so as to embrace a set of values and attributes’, resulting in a ‘powerfully held set of beliefs by the consumer.’²⁵⁶ Whereas trademarks have been thought to convey limited information about source or quality,²⁵⁷ brands have been understood to have the capacity for a more expansive function.²⁵⁸ A WIPO report notes that “[b]rands are not merely viewed as instruments for differentiation.”²⁵⁹ Brands are not a proxy for information, quality, or past experiences. Instead, brands convey associations, reputation, emotions, psychology, and image.²⁶⁰ A brand is said to be a company’s most valuable asset.²⁶¹ It is a means to dominate a market segment.²⁶²

Brands are manufactured desire. They are products of advertising. They are not the label on a product, but the product itself.²⁶³ As Professor Rochelle Dreyfuss observed, “ideograms that once functioned solely as signals denoting the source, origin, and quality of goods, have become products in their own right, valued as indicators of the status, preferences,

Biasing Brands, *supra* note 59, at 1259 (“[T]he marketing and psychology literatures have arrived at somewhat different conclusions concerning the nature of consumer decision-making than those offered by . . . the search-costs model.”); Rochelle C. Dreyfuss, *Expressive Genericity: Trademarks as Language in the Pepsi Generation*, 65 NOTRE DAME L. REV. 397, 398 (1990) (“Trademark law has not kept pace with trademark practice.”).

256. PETER URWIN, VALERIYA KARUK, PHILIP HEDGES & FRANK AUTON, VALUING BRANDS IN THE UK ECONOMY 9 (2008) (footnotes omitted).

257. *See supra* section II.B; 3 MCCARTHY, *supra* note 22, § 18:55.

258. *See generally* Lane, *supra* note 168; AAKER, *supra* note 254. Because the concept of brand includes “brand experience,” and is much larger than any one trademark, firms may use utility patents, design patents, copyrights, and trade secrets, in addition to trademarks, to protect their business model. *See* 2013 WORLD INTELLECTUAL PROPERTY REPORT, *supra* note 252, at 11, 13, 68–69.

259. 2013 WORLD INTELLECTUAL PROPERTY REPORT, *supra* note 252, at 22.

260. *See* THEY MAKE DESIGN, *Revealing the Psychology of Branding: Strategies for Success*, MEDIUM, <https://medium.com/theymakedesign/psychology-of-branding-5f64ba76f8> [<https://perma.cc/H973-HH82>]; *see also* Amanda Hess, *What Happens when People and Companies Are Both Just ‘Brands’?*, N.Y. TIMES MAG. (May 1, 2018) <https://www.nytimes.com/2018/05/01/magazine/what-happens-when-people-and-companies-are-both-just-brands.html> (last visited Nov. 18, 2023).

261. *See* AAKER, *supra* note 254, at 14 (“For many businesses the brand name and what it represents are its most important asset . . .”); *see also* MARK SHERRINGTON, ADDED VALUE: THE ALCHEMY OF BRAND-LED GROWTH 70 (2003); R.S.N. PILLAI, MARKETING MANAGEMENT 239 (2010).

262. For example, Glynn Lunney observed that branded products constitute their own separate product market. Lunney, *supra* note 29, at 430.

263. *See* Litman, *supra* note 244, at 1727–28.

and aspirations of those who use them.”²⁶⁴ They are Veblen goods²⁶⁵ that are conspicuously consumed.

When consumers desire the brand as the product in and of itself, the trademark ceases to equate to information about source of origin. That is, the material reason for the purchase is to own the brand and not for the unobservable information about the product that the trademark conveys.

The function of these marks may still be informational, but not to inform the purchaser; they are to inform the purchaser’s peers. Consumers are still attracted to brands, but for their status value and far less so for their informational value. Today, brand value has begun to supplant trademark value, but trademark law has not been adjusted accordingly.

Trademark law has not yet acknowledged this profound yet subtle shift from trademark to brand. It continues to protect trademarks for the informational value they theoretically offer while trademark owners, marketers, and consumers are increasingly only interested in a brand’s reputational value.

B. *Shopping in the Metaverse*

The shift from trademark to brand was prescient for trademark owners. In the virtual world, consumers will not need trademarks to inform them of a good’s features or qualities. Virtual goods are not made from better or worse quality leather, they are not hand- or machine-made, and they do not originate in Italy or China. But a handbag indistinguishable in quality, characteristics, or source may be expensive or inexpensive because it may be a Gucci handbag or a ZGGCD handbag. In the virtual world, brands’ status-communicating function is enhanced while trademarks’ information function is diminished.

The introduction of synthetic media—AI-generated content—has brought about a paradigm shift in marketing.²⁶⁶ Virtual spaces present marketers with strategic moments to expose consumers to visual prompts. There are increasing opportunities to shop within virtual and augmented reality, video games, and the metaverse.²⁶⁷ Marketers are attracted to these spaces because not only are their target audiences spending larger

264. Dreyfuss, *supra* note 255, at 397.

265. See generally James Chen, *Veblen Good: Definition, Examples, Difference from Giffen Good*, INVESTOPEDIA (June 29, 2023), <https://www.investopedia.com/terms/v/veblen-good.asp> [<https://perma.cc/MAC5-GGGA>] (“A Veblen good is a good for which demand increases as the price increases due to its exclusive nature and appeal as a status symbol.”).

266. Venkataramakrishnan, *supra* note 219.

267. See, e.g., Murphy, *supra* note 192 (“[Meta] intends to cash in on its virtual world, with hyper-targeted advertising and sponsored content that mirrors its existing \$85bn-a-year ad-based business model.”).

amounts of time there,²⁶⁸ but they also pay better attention to visual stimuli in these spaces. One arena in which this has become common is in video games where players are exposed to in-game product placement.²⁶⁹ Here marketers merge advertising with content with built-in feedback in real-time.

NFTs and the metaverse are the new frontier for brands. Some commentators suggest that the market for branded products in the metaverse could rival the market for branded products in the analog world.²⁷⁰ Already, firms as diverse as Nike, McDonald's, Victoria's Secret, and Walmart, have filed trademark applications for registration for a range of *virtual* products.²⁷¹ Moreover, trademark disputes have begun even before the metaverse is fully ready for commerce.²⁷² Marketers and brand owners are already eyeing NFTs and the metaverse as “an exciting way for brands to interact with their consumers.”²⁷³ This interaction of brands with consumers in a virtual space is something not contemplated by search costs theory. Consumers will buy virtual Nike shoes for their avatar neither for the quality of those virtual shoes nor for their origin.

The metaverse will be for trademarks the equivalent of stage actors breaking the fourth wall. Here it is impossible to think of the trademark as having any informational value. In the metaverse, it is explicit that the products are the brands.²⁷⁴

268. PRING ET AL., *supra* note 198, at 42.

269. Egliston, *supra* note 229 (“Increasingly, video game companies exploit this data to capitalize user attention through targeted advertisements.”).

270. Bernard Marr, *How Luxury Brands Are Making Money in the Metaverse*, FORBES (Jan. 19, 2022, 1:56 AM EST), <https://www.forbes.com/sites/bernardmarr/2022/01/19/how-luxury-brands-are-making-money-in-the-metaverse/> [<https://perma.cc/MXD9-E49F>] (“Analysts at Morgan Stanley say the market for virtual luxury goods could be as large as \$50 billion by 2030.”).

271. Yadira Gonzalez, *Brands That Have Filed Metaverse Trademarks—and What It All Means*, ADAGE (Apr. 11, 2022), <https://adage.com/article/digital-marketing-ad-tech-news/brands-have-filed-metaverse-trademarks/2401886> (last visited Nov. 7, 2023).

272. See, e.g., Complaint, Nike, Inc. v. StockX LLC, No. 22-cv-983 (S.D.N.Y. Feb. 3, 2022); Complaint, Miramax, LLC v. Tarantino, No. 2:21-cv-08979 (C.D. Cal. Nov. 16, 2021); James Ellis, *Non-Fungible Olive Gardens NFT Delisted on OpenSea!*, NFT EVENING (Sept. 21, 2022), <https://nftevening.com/non-fungible-olive-gardens-nft-delisted-on-opensea/> [<https://perma.cc/T2ND-PSPG>] (discussing the attempts of Darden Restaurants, owner of the Olive Garden brand, to take down a collection of NFTs depicting real-world Olive Gardens that come with a set of tokens called “Breadsticks”). Hermès won one of the first lawsuits to address the intersection of trademarks and NFTs. *Hermès v. Rothschild: A Timeline of Developments in a Case Over Trademarks, NFTs*, THE FASHION LAW (Nov. 20, 2023), <https://www.thefashionlaw.com/hermes-v-rothschild-a-timeline-of-developments-in-a-case-over-trademarks-nfts/> [<https://perma.cc/8ZD6-Z965>] (“[T]he case has proven to be a closely-watched matter due to its status as one of the first lawsuits to center on the intersection of trademarks and NFTs and its focus on key questions, including the extent to which ‘real’ world trademark rights extend to the virtual world.”).

273. Complaint at 1, Nike, Inc. v. StockX LLC, No. 22-cv-983 (S.D.N.Y. Feb. 3, 2022).

274. Desai, *supra* note 250.

C. *Trademarks Need a New Rationale*

A trademark's source identifying function can be broken down into two interrelated functions. The most basic is the referential function. Trademarks "identify and distinguish"²⁷⁵ a particular product from others. This function enables consumers to easily and efficiently spot the product they seek because the trademark sets it apart from other products. Although this function reduces the time consumers spend scanning for the desired product, the more time-consuming step is researching which product they desire to purchase. Enter the information transmission function. Trademarks enable consumers to take whatever they know about a product (from experience, recommendations, reviews, advertising) and ascribe that information to a product that bears the associated mark. The mark becomes a stand in for research about the quality and characteristics they seek in a product. It is the reduction of the cost of product research that makes trademarks worthy of legal protection.

Should the law protect trademarks to the same extent if they now perform only the referential function? Trademarks still serve as a referent for the product about which consumers can research and get information. But trademarks' information transmission function is waning as consumers now have more accurate and efficient means of conducting product research.

There is a further dissonance between how trademarks actually function and what function the law purports to protect. The information that trademarks purport to transmit has become unhinged from any particular goods. Rather than information about quality and characteristics, it has become free-floating information about associations that ride with the mark. The law then is protecting the communication of a merchant's message. This protection has not yet been justified.

Trademark law is replete with references to the consumer and doctrines that are seemingly built on consumers' experiences. The story we tell ourselves is that trademark law is about protecting consumers.²⁷⁶ However, trademark law is not a form of consumer protection law and

275. This phrase appears in the Lanham Act's definition of a trademark: "any word, name, symbol, or device or any combination thereof . . . used by a person . . . to identify and distinguish his or her goods . . . from those manufactured or sold by others." 15 U.S.C. § 1127.

276. MCCARTHY, *supra* note 22, § 2:1 ("The interest of the public in not being deceived has been called the basic policy [of trademark law]."); Bone, *Hunting Goodwill*, *supra* note 29, at 549 ("The core of trademark law . . . views trademarks as devices for communicating information to the market and sees the goal of trademark law as preventing others from using similar marks to deceive or confuse consumers."). *But see* McKenna, *Normative Foundations*, *supra* note 249, at 1916 ("Critics cannot continue simply to claim that modern law is illegitimate because it does not seek to protect consumers. Because it never really did.").

never has been.²⁷⁷ As has been shown, consumers are not benefitted by trademarks as much as the law's rationale holds, and in some cases, they are likely harmed by the misinformation trademarks provide.

Before the informational function of trademarks became doctrinaire, trademark theorists frankly engaged with the policy justifications for trademarks. Just after the passage of the Lanham Act, Professor Ralph Brown warned against broad unexamined justifications of protection. Openly questioning whether trademarks and advertising serve the public interest, Brown concluded that they did only when they provided consumers useful information about products.²⁷⁸ Informing, however, was different from persuading, which disserves the public and is not worthy of protection.²⁷⁹ Brown's approach has salience today. When trademarks do not play an information transmission function, but instead function as persuasive advertising, the reason for protecting them disappears.

Should trademarks be protected when they fail to inform? Perhaps, but not without a rationale that matches the facts on the ground. Professor Felix Cohen, another early trademark theorist, argued that trademark law had abandoned its role of protecting consumers in favor of protecting property interests and had failed to offer a theory as to why protecting these interests is socially valuable.²⁸⁰ The Chicago School's search costs theory might be seen as a response to Cohen's attack on the social utility of trademarks.²⁸¹ Cohen's critique, however, was larger. His insight was that legal concepts disconnected from empirical observations are "transcendental nonsense."²⁸² Search costs theory has long suffered from a disconnect with facts on the ground, but our increasingly algorithmic world has now made this apparent.

Trademark scholar Frank Schechter's contributions can be read as frustration that the account of trademarks as information was incomplete. Schechter's famous 1927 Harvard Law Review article argued for

277. McKenna, *Normative Foundations*, *supra* note 249, at 1841 (arguing that trademark law was not originally intended to protect the consumer, but instead sought to protect producers from "illegitimate diversions of their trade by competitors").

278. Ralph S. Brown, Jr., *Advertising and the Public Interest: Legal Protection of Trade Symbols*, 57 *YALE L.J.* 1165, 1168, 1185, 1187 (1948) (arguing that protecting trademarks' "informational value" is legitimate).

279. *Id.* at 1168 ("[I]mparting information is the only useful function of advertising."); *id.* at 1183–84, 1190.

280. See Felix S. Cohen, *Transcendental Nonsense & the Functional Approach*, 35 *COLUM. L. REV.* 809, 814 (1935) ("There was once a theory that the law of trade marks and tradenames was an attempt to protect the consumer against the 'passing off' of inferior goods under misleading labels.")

281. *Peaceable Planet, Inc. v. Ty, Inc.*, 362 F.3d 986, 990 (7th Cir. 2004) (noting that one general policy "rationale" of trademark law has come to be to "promote competition and consumer welfare").

282. Cohen, *supra* note 280, at 815 ("Courts and scholars, therefore, have taken refuge in a vicious circle to which no obviously extra-legal facts can gain admittance.")

expanded trademark rights to reflect an expanded understanding of how trademarks actually functioned in the marketplace.²⁸³ Schechter maintained that trademarks are better understood as “creating a desire.”²⁸⁴ Schechter may have been more accurate than the prevailing account, but as Brown and Cohen later demonstrated, such a function was difficult to defend policy-wise.²⁸⁵ In contrast, Schechter’s contemporary, trademark attorney Edward S. Rogers, seems to have appreciated the plasticity of the source of origin confusion model.²⁸⁶ As the main drafter of the Lanham Act, Rogers’s view prevailed. Although search costs theory is an unequivocal rejection of Schechter’s account, today, we have nevertheless come to protect trademark’s selling power.

The remarkable staying power of the search costs account of trademarks, however, has provided judges cover to create, protect, and reinforce monopoly powers while giving lip service to consumer protection. Instead, courts should have to acknowledge the protection they are granting and in whose interest it is. Brand owners should consciously rethink the rationale of protecting brands as opposed to trademarks. That entails protecting brands qua brands, not brands as a proxy for other information. Courts should understand that protection aids only the brand owner, not the consumer, and not competition. Perhaps there is a balance where the law can appropriately reward the brand owner for their investment in branding, while safeguarding competition. Search costs theory, however, provides no room for such a balance.

Retail has changed more in the past two decades than it has over the entire course of trademark history and even bigger changes are just around the corner. We can think of this as the third major period of retail, the first being the era of the local dry goods store where consumers did not interact

283. See generally Schechter, *supra* note 26.

284. *Id.* at 819. In a similar vein, in 1942, Justice Frankfurter somewhat infamously described trademark protection as “the law’s recognition of the psychological function of symbols” and referred to trademarks as “a merchandising short-cut” meant “to impregnate the atmosphere of the market with the drawing power of a congenial symbol.” *Mishawaka Rubber & Woolen Mfg. Co. v. S.S. Kresge Co.*, 316 U.S. 203, 205 (1942).

285. See Brown, *supra* note 278, at 1166. Judge Jerome Frank also recognized that protecting trademarks in the name of protective consumers posed the danger of creating legal monopolies. *Triangle Publ’ns, Inc. v. Rohrlich*, 167 F.2d 969, 980 n.13 (2d Cir. 1948) (Frank, J., dissenting) (“The trade-name doctrine . . . enables one to acquire a vested interest in a demand ‘spuriously’ stimulated through ‘the art of advertising’ by ‘the power of reiterated suggestion’ which creates stubborn habits. . . . Should the courts actively lend their aid to the making of profits derived from the building of such habits, if and whenever those stubborn habits so dominate buyers that they pay more for a product than for an equally good competing product? . . . [I]ts basic virtue is generally regarded as consisting of its benefits to consumers.”).

286. See Jessica Litman, *Edward S. Rogers, the Lanham Act, and the Common Law*, in ROBERT G. BONE & LIONEL BENTLEY, *RESEARCH HANDBOOK ON THE HISTORY OF TRADEMARK LAW* 1–43 (forthcoming 2023) (on file with author).

with trademarks at all, and the second being the marketplace that search costs theory described. The second period ushered in our modern trademark law. The change in the marketplace in the third period is as fundamentally different from the second period as the second was from the first. Such a major change in the marketplace should occasion a concomitant change in trademark theory.

Trademarks may still sometimes assist consumers in purchasing decisions, but they may just as likely lead to less optimal results, slow down search, distract, or even misinform. Today—and likely much more so tomorrow—consumers face a different shopping experience in which trademarks are displaced by more accurate and efficient tools. Consumers’ reliance on trademarks for information about the source or quality of products will progressively fade.

Trademark theory needs to acknowledge that trademarks are doing less work and doing it less often. Instead of protecting trademarks qua information, the law is protecting brands. Brands are valuable properties because they function in the marketplace as the product itself. Presently, brands enjoy robust protection because of unexamined application of search costs theory and not because of a considered determination that they are deserving of legal protection.

Trademark law needs to recognize when search costs are and are not implicated and measure protection accordingly. In a number of concurring and dissenting opinions in the years immediately preceding and following the enactment of the Lanham Act, Second Circuit Court of Appeals Judge Frank directly challenged the informational value of trademarks to consumers stating that “the conventional assumption that trade-name protection importantly adds, in direct fashion, to consumers’ economic welfare, has not as yet been proved to be true in fact.”²⁸⁷ Frank pointedly asked: if trademark law’s “exceptions to the presumption in favor of competition are of no direct use to consumers, do they serve a sufficiently important social interest to justify their existence?”²⁸⁸

This Article contends that the support for the economic welfare or social utility claim is now at its weakest point in the history of trademark law. It is time to recalibrate trademark law’s normative foundations.

287. *Triangle Publ’ns, Inc.*, 167 F.2d at 980 n.13 (Frank, J., dissenting); see also *Standard Brands v. Smidler*, 151 F.2d 34, 37 (2d Cir. 1945) (Frank, J., concurring); *LaTouraine Coffee Co., Inc. v. Lorraine Coffee Co., Inc.*, 157 F.2d 115, 124–25 (2d Cir. 1946) (Frank, J., dissenting).

288. *Triangle Publ’ns, Inc.*, 167 F.2d at 980 n.13.

CONCLUSION

The claim that trademarks transmit useful information and that consumers rely on this function to efficiently navigate the marketplace has been undermined by the age of algorithms. The normative rationale for protecting trademarks is rapidly slipping away. Trademarks may still function to distinguish products on occasion, but changes in technology indicate that these occasions are dwindling.

These technologies have completely transformed the very nature of how sellers sell. Search costs theory assumes the unavailability of alternative means to address the asymmetry of information between producers and consumers. Now, however, the marketplace utilizes an arsenal of digital technologies to enable consumers to efficiently connect with products that meet their desires. The need to rely on trademarks to convey information about products is therefore greatly diminished. These technological alternatives dramatically change the value of trademarks for consumers and so should change the cost benefit analysis of protecting them.

In the early days of retail in America, consumers would not browse the store shelves to find the products they wanted to buy. Instead, their local shopkeeper would fetch them products that corresponded to their stated needs. In this way, we have come full circle where consumers today again delegate purchasing decisions. As a result of unleashing AI on consumers' harvested personal data, firms know the needs *and* desires of consumers, maybe even before consumers do. In this scenario, as in the earliest local shops, search costs theory is inapt.

Even in the twentieth century marketplace to which search costs theory was addressed, it failed to acknowledge how trademarks often functioned to obfuscate rather than reveal source information. Increasingly, trademark law has served to protect property interests in the name of protecting consumers. Search costs theory shields the work trademark law is actually doing to protect brands as the products themselves. That trademarks now play an increasingly different role in society should prompt a reconsideration of whether trademark protection is justified in particular cases.

TRADEMARK CONFUSION REVEALED: AN EMPIRICAL ANALYSIS

DARYL LIM*

The likelihood of confusion standard defines the scope of trademark infringement. Likelihood of confusion examines whether there is a substantial risk that consumers will be confused as to the source, identity, sponsorship, or origin of the defendants' goods or services. This Article presents a contemporary empirical analysis of the various factors and how they interact. Conventional wisdom teaches us that courts should comprehensively traverse each factor and that likelihood of confusion cases generally require jury determination. However, the data reveals that neither is true. Instead, courts provide early off-ramps to litigants by "economizing," and analyzing only a handful of factors or by "folding" factors within each other. The findings also reveal (1) which forums are pro-defendant and which are pro-plaintiff; (2) the impact of rivalry and fair use on outcomes; and (3) an apparent Ninth Circuit dominance.

What constitutes "confusion" remains highly subjective and difficult to evaluate. Proxies like intent, survey evidence, mark strength, and consumer sophistication fail to incorporate real-world purchasing conditions or are better considered within omnibus factors. In contrast, actual confusion, mark similarity, and competitive proximity provide judges with a potent trio of factors to guide the infringement inquiry. Together with safe harbors for descriptive and expressive uses, these rules of thumb enable courts to resolve trademark disputes more coherently, consistently, and expeditiously. This Article concludes with a blueprint of how these rules of thumb complement artificial intelligence systems

* Professor of Law & Director, Center for Intellectual Property, Information and Privacy Law, *University of Illinois Chicago School of Law*. I thank Professor Y. Samuel Wang from the Department of Statistics and Data Sciences at *Cornell University* for his valuable advice on the statistical aspects of this Article. My sincere thanks to Margaret Smiley Chavez, Steve Fisher, Annemarie Gregoire, Sarah Hampton, Nicole Robinson, Kelly Welsh, and their colleagues at the *American University Law Review* who contributed to this Article. Their courtesy, professionalism, careful editorial work, and thoughtful comments are exemplary.

and how those systems can use empirical studies as training data to inform future likelihood of confusion analyses.

TABLE OF CONTENTS

Introduction	1286
I. Studying Confusion	1290
A. Methodology	1291
B. Blends, Triggers, and <i>Polaroid</i> Factors	1296
1. Blends and triggers	1297
2. The <i>Polaroid</i> factors	1300
C. Missing the Point on Consumer Confusion	1306
1. Intent	1306
2. Surveys.....	1312
3. Mark strength	1316
4. Consumer sophistication	1319
D. Coherence-Based Reasoning.....	1320
II. Rules of Thumb.....	1324
A. Actual Confusion.....	1324
B. Mark Similarity	1328
C. Competitive Proximity	1332
D. Summing It Up	1338
E. A Word on Fair Use.....	1340
III. Observations and Implications.....	1345
A. Factor Folding	1345
B. Early Off-Ramps	1347
C. Deploying Artificial Intelligence	1353
1. Conception	1354
2. Execution.....	1355
3. Limitations.....	1358
a. “Garbage-in, garbage-out”	1358
b. Biases	1359
c. Contextualizing the purchasing conditions.....	1362
d. Coding challenges	1364
Conclusion.....	1365

INTRODUCTION

Consumers rely on a consistent commercial lexicon to reduce mental costs associated with purchasing decisions and in turn commercial enterprises gain an incentive to invest in quality products

and services.¹ Businesses imbue words, symbols, scents, and sounds with information about their goods and services.² In turn, consumers rely on this imbued information to navigate their decisions by making snap judgments about the price and quality of products or services they purchase.³ Thus, electric vehicle enthusiasts may seek out Tesla's stylized "T" comprising a stator and rotor, and connoisseurs of Chick-fil-A's chicken sandwiches will scout for its distinctive red-on-white text.

When trademark owners assert their rights, courts apply the likelihood of confusion standard, which seeks to determine whether the defendants' use of a similar mark is likely to cause consumer confusion. The likelihood of confusion standard is the linchpin of trademark infringement.⁴

Unfortunately, what constitutes "confusion" remains highly subjective and difficult to evaluate.⁵ Additionally, the likelihood of confusion standard remains poorly theorized and judges applying the standard often fail to adequately explain their analyses in their opinions in a way that future courts can easily apply and replicate.⁶

When defendants counterfeit the trademark outright, liability is clear. Literal infringement has occurred. However, trademarks protect their owners beyond literal infringement like patents and copyrights.⁷ Nonliteral infringement can occur when, for example, there is a colorable difference in the marks. This exposes parties to uncomfortably

1. *Qualitex Co. v. Jacobson Prods. Co.*, 514 U.S. 159, 163–64 (1995) (“[T]rademark law . . . reduces the customer’s costs of shopping and making purchasing decisions, . . . for it quickly and easily assures a potential customer that this item—the item with this mark—is made by the same producer as other similarly marked items that he or she liked (or disliked) in the past.”).

2. See 1 J. THOMAS MCCARTHY, MCCARTHY ON TRADEMARKS AND UNFAIR COMPETITION § 3, at 1 (5th ed., 2010).

3. *Qualitex Co.*, 514 U.S. at 163–64.

4. See *infra*, Part I. The Lanham Act prohibits the use of a registered mark in a manner “likely to cause confusion,” 15 U.S.C. § 1114(1)(a), as well as the use of any term or name in a manner “likely to cause confusion” about the affiliation of the user with another person. *Id.* § 1125(a)(1); see also *Mil. Ord. of Purple Heart Serv. Found., Inc. v. Mil. Ord. of Purple Heart of U.S., Inc.*, 852 F. App’x 6, 9 (D.C. Cir. 2021).

5. See *infra* Part I.

6. See *infra* Part II; see also Barton Beebe, *An Empirical Study of the Multifactor Tests for Trademark Infringement*, 94 CALIF. L. REV. 1581, 1582 (2006) (“Its current condition is Babelian.”).

7. See generally Daryl Lim, *Judging Equivalents*, 36 SANTA CLARA HIGH TECH. L.J. 223 (2020) (tracing the origins of the doctrine of equivalents and explaining the rationale behind the doctrine—to protect intellectual property owners from infringers seeking to evade liability by making insubstantial changes) [hereinafter Lim, *Judging Equivalents*].

uncertain waters.⁸ Patent law has claims to give notice of a patent's metes and bound.⁹ Neither trademark nor copyright law has claims, leaving courts unclear on operationalizing technical similarity or market substitution considerations.¹⁰

In a negative feedback loop, indeterminacy over the likelihood of confusion standard muddies trademark law's focal point and scope, while polluting adjacent disciplines like copyright and patent law.¹¹ Proper notice about the existence and scope of legal rights is critical to any property system, but especially trademark rights, because trademarks last indefinitely, meaning their owners obtain a timeless monopoly without the same limitations and threshold requirements placed on patent and copyright holders.¹² A patchwork of inconsistent results destabilizes the system for everyone.¹³ Indeterminacy has many negative impacts, including causing negotiations to break down, which harms *both* brand owners and potential licensees, and acting as a drag on dispute resolution, compliance, and social equity.¹⁴ Indeterminacy also acts as a drag on dispute resolution, compliance, and social equity.¹⁵ The rational response must be a call for clarity in the law.

The likelihood of confusion standard examines whether consumers will likely be confused as to the source, identity, sponsorship, or origin of the goods and requires "a substantial likelihood that the public will

8. Michael Grynberg, *Thick Marks, Thin Marks*, 67 CASE W. RESV. L. REV. 13, 15 (2016) ("Many open questions in modern trademark law concern which parts of the range belong under the trademark holder's control.")

9. 35 U.S.C. § 112(b) (requiring patentees to include in their patent "one or more claims particularly pointing out and distinctly claiming the subject matter which the inventor . . . regards as the invention").

10. Jeanne C. Fromer & Mark A. Lemley, *The Audience in Intellectual Property Infringement*, 112 MICH. L. REV. 1251, 1296–99 (2014).

11. See Robert G. Bone, *Notice Failure and Defenses in Trademark Law*, 96 B.U. L. REV. 1245, 1255 (2016) ("[W]hat makes the scope of rights so uncertain is the vagueness of the likelihood-of-confusion test ("LOC test") for infringement.") [hereinafter Bone, *Notice Failure and Defenses*]; Amy Adler & Jeanne C. Fromer, *Taking Intellectual Property into Their Own Hands*, 107 CALIF. L. REV. 1455, 1523 (2019) ("Trademark law is similarly complex and unpredictable with regard to important doctrines.")

12. See William M. Landes & Richard A. Posner, *Trademark Law: An Economic Perspective*, 30 J.L. & ECON. 265, 287 (1987) ("The lack of a fixed term for trademarks is one of the striking differences between trademarks, on the one hand, and copyrights and patents, on the other."). See generally 15 U.S.C. § 1114 (discussing infringement and remedies).

13. Thomas H. Watson, *Pay Per Click: Keyword Advertising and the Search for Limitations of Online Trademark Infringement Liability*, 2 CASE W. RESV. J.L. TECH. & INTERNET 101, 122 (2011).

14. See, e.g., Bone, *supra* note 11, at 1258.

15. Daryl Lim, *AI, IP, Algorithms, and Inequality*, SMU L. REV. (forthcoming 2022).

be confused.”¹⁶ Each circuit court has enumerated factors relevant in analyzing likelihood of confusion. Courts use proxies for consumer confusion like intent, survey evidence, mark strength, and consumer sophistication.¹⁷ However, these either fail to incorporate real-world purchasing conditions as a doctrinal matter or are better considered part of a streamlined likelihood of confusion test.¹⁸ The trio of actual confusion, mark similarity, and competitive proximity provides judges a small but potent cluster of factors.¹⁹ Together with safe harbors for descriptive and expressive uses, these enable a more coherent, consistent, and expedient resolution of trademark disputes.

This Article presents a contemporary empirical analysis of each likelihood of confusion factor and how they interact with one another. Conventional wisdom teaches us that courts should comprehensively traverse each factor and that likelihood of confusion cases generally require jury determination.²⁰ The data reveals that neither is true. Instead, courts provide early off-ramps to litigants by “economizing” and applying just a handful of factors or by “folding” factors into each other in grouped layers.²¹ The findings also reveal the Ninth Circuit’s dominance, pro-defendant and pro-plaintiff fora, and the impact of rivalry and fair use on case outcomes.²² This Article concludes with a blueprint of how artificial intelligence (AI) systems can use empirical studies as training data to help stakeholders make and predict confusion analyses.²³

Part I introduces this empirical study’s methodology, goals, and limitations before elaborating on this Article’s doctrinal and policy impetus. The discussion charts how blending technical trademarks and trade names along with the expansion of triggers ensnares defendants in trademark liability. It then shifts to make critical observations gleaned from the data, including the impact of rivalry on modern case outcomes, the dominance of the Ninth Circuit in federal trademark litigation, and the most pro-defendant and pro-plaintiff circuits today.

16. *Fisher Stoves, Inc. v. All Nighter Stove Works, Inc.*, 626 F.2d 193, 194 (1st Cir. 1980).

17. *See infra* Section I.C.

18. *See generally* 4 MCCARTHY, *supra* note 2, § 23:1 (criticizing the “amorphous and indefinite” nature of the multi-factor likelihood of confusion test for creating legal unpredictability, increasing litigation costs, and chilling socially valuable uses).

19. *See infra* Part II.

20. *See infra* Part II.

21. *See infra* Part III.

22. *See infra* Part I.

23. *See infra* Part III.

Part I also presents a detailed doctrinal and empirical analysis of prominent features in the likelihood of confusion analysis—the intent of litigating parties, consumer surveys, mark strength, and consumer sophistication—and explains why each factor leads courts tasked with ascertaining consumer confusion astray. Finally, the Part introduces coherence-based reasoning and argues that a more compact nucleus of factors would better serve courts analyzing likelihood of confusion.

Part II explains why actual confusion, mark similarity, and competitive proximity offer courts that compact troika in simplifying likelihood of confusion analysis. The data reveals how most courts rely on these three factors while either paying lip service to or completely ignoring the other factors. It also shows appellate courts are complicit in this “wink-and-nod” practice, affirming lower courts in over 80% of cases on appeal. Finally, Part II explains why this trio of factors plus the fair use safe harbors of descriptive and expressive uses should form trademark law’s rules of thumb for infringement.

Part III examines the implications of the empirical study on trademark doctrine and practice. First, it observes that while the likelihood of confusion factors may present themselves as discrete categories, the dataset reveals that courts do not regard them as such. Instead, courts combine factors and analyze them both creatively and rationally. Second, courts rely on a small number of factors to economize their decisions and give parties an early off-ramp. Third, Part III explains how the empirical analysis provides a blueprint for algorithmic adjudication using AI, taking the reader from conception to execution while identifying and addressing its limitations.

I. STUDYING CONFUSION

Over the years, the jurisprudential roots of trademark law became unruly and tangled. Unfair competition intermingled with consumer protection as the Lanham Act of 1946²⁴ (“the Act”) blended trade names and technical trademarks.²⁵ A later legislative revision untied likelihood of confusion from source confusion, and courts introduced idiosyncratic rules of affiliation and sponsorship as triggers for consumer confusion.²⁶

Within the likelihood of confusion tests, factors such as defendants’ intent, survey evidence, and consumer sophistication provided a convenient but misguided attempt to get a handle on the arduous task

24. 15 U.S.C. § 1051–1127.

25. *See infra* Section I.B.

26. *See infra* Section I.B.

of determining what trademark infringement had become.²⁷ Judges often resorted to coherence-based reasoning. Once a judge determined the satisfaction of the selected factors, the judge would decide that all the other factors were present.²⁸ Though it made their work easier, it muddied the waters for everyone else. This Part presents an empirical analysis of these issues and their implications. It begins by discussing the empirical methodology.

A. Methodology

This empirical study draws upon the well-developed method of case content analysis.²⁹ The method systematically dissects a sample of judicial opinions to record consistent features, draw inferences, and uncover trends.³⁰ This social science approach to the law complements and augments traditional legal analysis.³¹ As a testament to its outsized contribution to the literature, case content analysis generates an average of seventy-seven citations per article in a milieu where 40% of law review articles receive no citations at all.³²

The value of case content analysis lies in uncovering patterns in judging. Judges routinely rely on a remarkably limited number of factors in reaching their conclusions.³³ Instead, they employ

27. See *infra* Section I.B.

28. See *infra* Section I.D.

29. Mark A. Hall & Ronald F. Wright, *Systematic Content Analysis of Judicial Opinions*, 96 CALIF. L. REV. 63, 73 (2008) (“Content analysis has proven useful for studying a broad range of legal subject areas.”); *id.* (listing “areas as far-ranging as administrative law, constitutional law, corporate and securities law, criminal law and procedure, contracts, employment discrimination, health law, and torts”).

30. See KLAUS KRIPPENDORFF, *CONTENT ANALYSIS: AN INTRODUCTION TO ITS METHODOLOGY* 18 (2d ed., 2004) (defining content analysis as “a research technique for making replicable and valid inferences from texts (or other meaningful matter) to the contexts of their use”).

31. Hall & Wright, *supra* note 29, at 74; see also *id.* at 65 (“The method also helps a researcher to sort out the interaction of multiple factors that bear on an outcome in the legal system.”); *id.* at 78 (“Its strength is to provide an objective understanding of a large number of decisions where each decision has roughly the same value.”); Alan L. Tyree, *Fact Content Analysis of Case Law: Methods and Limitations*, 22 JURIMETRICS J. 1, 23 (1981) (explaining that these methods have “considerable power for the discovery of anomalies which may escape the naked eye”).

32. Hall & Wright, *supra* note 29, at 74 (“[C]ontent analysis projects appear somewhat more likely to generate discussion and citation than law review articles more generally.”).

33. See, e.g., Beebe, *supra* note 6, at 1602 (“The data collected for this study support the general hypothesis that decision makers, even when making complex decisions, consider only a small number of factors and the more specific hypothesis that, in doing so, decision makers use a core attributes heuristic.”).

heuristics—such as the likelihood of confusion factors—to cut through what would otherwise be a morass of information that could paralyze decision-making entirely.³⁴ This approach, however, makes it difficult to draw broader conclusions to inform future cases coherently.³⁵ Scholars employ case content analysis to parse through court decisions and study how judges and juries apply rules to facts in practice to address this limitation.³⁶ Thus, while the interpretive method evaluates legal principles, case content analysis “combines the analytical skills of the lawyer with the power of science that comes from articulated and replicable methods of reading and counting cases.”³⁷ In so doing, case content analysis yields useful information that moves the discussion toward a greater understanding of the bigger policy questions and helps uncover areas for further research.³⁸ It also avoids selection bias issues, which hamper the representativeness of other methods.³⁹

This Article relied on an expansive pool of 188 cases covering nearly 5,000 datapoints based on a Westlaw search for all trademark infringement cases discussing likelihood of confusion over five years between September 30, 2016, and October 1, 2021.⁴⁰ The study

34. See, e.g., Daryl Lim, *Retooling the Patent-Antitrust Intersection: Insights from Behavioral Economics*, 69 BAYLOR L. REV. 124 (2017).

35. See, e.g., Daryl Lim, *The (Unnoticed) Revitalization of the Doctrine of Equivalents*, ST. JOHN'S L. REV. (forthcoming 2022) [hereinafter Lim, *Doctrine of Equivalents*].

36. See, e.g., Lee Petherbridge, *On the Decline of the Doctrine of Equivalents*, 31 CARDOZO L. REV. 1371, 1380 (2010) [hereinafter Petherbridge, *Decline of the Doctrine of Equivalents*] (“Content analysis is capable of helping scholars verify, analyze, or refute empirical claims about case law, and it is to that purpose the approach is put in this study.”). For earlier studies where I employed a similar methodology, see Lee Petherbridge, Jason Rantanen, & Ali Majibi, *The Federal Circuit and Inequitable Conduct: An Empirical Assessment*, 84 S. CALIF. L. REV. 1293, 1304 (2011); Hall & Wright, *supra* note 29, at 77; DARYL LIM, PATENT MISUSE & ANTITRUST: EMPIRICAL, DOCTRINAL & POLICY PERSPECTIVES 8–9 (2013); Lim, *Judging Equivalents*, *supra* note 7.

37. Hall & Wright, *supra* note 29, at 100.

38. Karen A. Jordan, *Empirical Studies of Judicial Decisions Serve an Important Role in the Cumulative Process of Policy Making*, 31 IND. L. REV. 81, 88 (1998); see also Lon L. Fuller, *An Afterward: Science and the Judicial Process*, 79 HARV. L. REV. 1604, 1622 (1965) (“[P]ossible gain from researches of this kind lies in the realm of serendipity. A puzzling correlation that violates normal anticipations may set our minds going along new paths and yield unexpected insights.”).

39. Hall & Wright, *supra* note 29, at 102 (“All of these were universal samples restricted only by date, subject matter, jurisdiction, and/or source. In short, empirical researchers studying case law are usually able to avoid the selection bias issues that plague most other areas of social science.”).

40. For an example of another recent empirical study on trademarks that starting and ending during the calendar year, see Lisa Larrimore Ouellette, *The Google Shortcut to Trademark Law*, 102 CALIF. L. REV. 351, 373 (2014) (“New U.S. federal court decisions

omitted thirty-two cases captured by Westlaw's search results because they did not specifically discuss the likelihood of confusion factors.⁴¹ For accuracy, the dataset distinguishes between procedural wins (for instance, defeating a motion for summary judgment) and substantive wins on the merits (which result in a finding of infringement or non-infringement).⁴² This Article initially used Excel to hand-code the data before using IBM SPSS Statistics 28.0 to generate the graphs and crosstabs data.⁴³

The dataset of hand-coded cases included variables such as (1) the decision's date; (2) the judicial circuit; (3) whether a district or appellate court decided the case; (4) the parties' relationship as rivals (or not); (5) the procedural posture; (6) the type of mark; (7) the test employed; (8) whether the opinion discussed survey evidence; (9) which party a likelihood of confusion factor favored; (10) whether courts "folded" factors together; (11) case outcomes; and (12) whether the court discussed fair use.⁴⁴

Like all empirical studies, this one has its caveats. There are several well-recognized limitations to case content study databases. First, coding may result in incomplete or inaccurate coding, despite cross-coding and verification using a population sample.⁴⁵ Given that the focus is on features of written decisions, the data remains valid as long as it is recognized to refer to a specific population rather than a sample of all cases in all possible worlds.⁴⁶ Second, cases from legal databases such as Westlaw are known to underreport jury decisions.⁴⁷ To some

related to trademark or service mark distinctiveness or likelihood of confusion were tracked from December 2011 to November 2012."); *see also* John R. Allison & Mark A. Lemley, *The (Unnoticed) Demise of the Doctrine of Equivalents*, 59 STAN. L. REV. 955, 963 (2007) ("[W]e collected every district court and court of appeals decision on the doctrine of equivalents that appeared in Westlaw . . ."). The number of cases this Article employs compares favorably with the norm, which ranges from less than 100 cases to 300 cases. *See* Hall & Wright, *supra* note 29, at 102 ("Of these 114 universal samples, only 11 coded more than 1000 cases, and 21 coded from 500 to 1000. Twenty-six of these projects coded fewer than 100 cases (with 13 of these fewer than 51), and 39 coded between 100 and 300.").

41. *See* Daryl Lim, Trademark Case Statistics (Dec. 17, 2021) (on file with the American University Law Review).

42. *Id.*

43. *Id.*

44. *Id.*

45. *See, e.g.,* Lim, *Doctrine of Equivalents*, *supra* note 35.

46. *Id.*

47. Allison & Lemley, *supra* note 40, at 963–64 ("The universe of all decisions is of course different from the universe of those reported in Westlaw, and in particular our study underreports jury decisions. But our focus on written decisions (both reported

degree, comparing it to other studies that employ similar methods to control for that feature can mitigate the effects of the underreporting.

Courts have found that “simple counts and percents are sufficient to document” a claim about case law trends, challenge conventional wisdom, or suggest further study issues.⁴⁸ The case-counting method codes the entire population of relevant cases.⁴⁹ Statistics are unnecessary to prove that sample cases are representative of a larger population.⁵⁰ Because the outcome of each case—the dependent variable of interest—has five possible categories, the most appropriate regression model would be a multinomial regression.⁵¹ However, each of the seven relevant factors also has five possible categories. Hence, the sample size required for a multivariate regression that would simultaneously test each category’s effect of each relevant factor is much larger than that of the current dataset.⁵² This Article refrained from presenting regression results and instead presented descriptive statistics. Specifically, it describes the observed distribution of case outcomes when conditioned on one or two relevant factors. Although this does not attest to a specific level of statistical significance, these values are still instructive for the reader.

Third, analysis of judicial opinions has well-known limitations.⁵³ Statistics fail to account for extralegal factors influencing outcomes, such as the state of the case record on appeal and judicial deliberations discussed in the opinion.⁵⁴ In addition, litigants may consider the expertise and reputation of the district court judge in deciding whether to appeal, introducing selection bias effects into the appellate

and unreported) allows us to parse the grounds for decision and the reasoning of the opinions.”).

48. Hall & Wright, *supra* note 29, at 118.

49. *Id.*

50. *Id.*

51. ALAN AGRESTI, CATEGORICAL DATA ANALYSIS 293 (3d ed. 2013).

52. *Id.*

53. See R. Polk Wagner & Lee Petherbridge, *Is the Federal Circuit Succeeding? An Empirical Assessment of Judicial Performance*, 152 U. PA. L. REV. 1105, 1128–29 (2004) (discussing unobserved reasoning, strategic behavior, and selection bias).

54. Harry T. Edwards & Michael A. Livermore, *Pitfalls of Empirical Studies that Attempt to Understand the Factors Affecting Appellate Decisionmaking*, 58 DUKE L.J. 1895, 1899 (2009).

data.⁵⁵ Moreover, most cases settle, so decided cases are a nonrandom subset of all cases.⁵⁶

Fourth, case content analysis trades depth for breadth. The complexity of trademark litigation also makes it difficult to generalize even from a study covering hundreds of cases.⁵⁷ Numbers do not reflect judicial rhetoric or more subtle clues about a judicial opinion's precedential value.⁵⁸ Case coding documents what judges do rather than draw normative implications from the observations.⁵⁹

Fifth, parties are not randomly distributed throughout the judicial districts.⁶⁰ Some district courts may hear more cases that eventually settle. Other courts may hear more cases where the parties file based on domicile. District court judges are therefore not assigned a random sample of patent lawsuits since they are assigned cases from the judicial district where they sit.⁶¹ Circumstances such as a particular judge or jury may cause a case to settle where the same case before another judge or jury could proceed to an appeal.⁶² This Article focuses on how appellate and lower courts interpret precedent. Those interpretations are not uniform and can never be so.⁶³

55. David Schwartz, *Practice Makes Perfect? An Empirical Study of Claim Construction Reversal Rates in Patent Cases*, 107 MICH. L. REV. 223, 243 (2008).

56. Jay P. Kesan & Gwendolyn G. Ball, *How Are Patent Cases Resolved? An Empirical Examination of the Adjudication and Settlement of Patent Disputes*, 84 WASH. U. L. REV. 237, 273–74 (2006) (finding that between 65% and 68% of all patent cases filed in three particular years were resolved via settlement or a probable settlement).

57. See David L. Schwartz, *Explaining the Demise of the Doctrine of Equivalents*, 26 BERKELEY TECH. L.J. 1157, 1188 (2011) (“Because patent litigation as a whole is so complex, it is incredibly complicated to develop and test empirical models.”); Petherbridge, *Decline of the Doctrine of Equivalents*, *supra* note 36, at 1380 (noting biases inherent in this approach such as “unobserved reasoning, selection bias, and strategic behavior”).

58. Hall & Wright, *supra* note 29, at 87.

59. *Id.* at 97 (“Still, imperfect data must suffice because observing actual behaviors and gauging true attitudes would be impossible or cost-prohibitive. Similarly, even though judge-reported facts may not ‘purport to be the real facts,’ they are ‘near enough so that the savings in labor justifies the approximation.’”).

60. Kimberly A. Moore, *Forum Shopping in Patent Cases: Does Geographic Choice Affect Innovation?*, 79 N.C. L. REV. 889, 924–25 (2001).

61. Schwartz, *supra* note 55, at 242.

62. See *id.* at 242 n.119 (“[I]f [judges] have a really tough case, they can put tremendous pressure on the parties to settle so there won’t be an appealable order.”) (alteration in original) (quoting Judge Richard A. Posner)).

63. See Jeffrey J. Rachlinski, *Does Empirical Evidence on the Civil Justice System Produce or Resolve Conflict?*, 65 DEPAUL L. REV. 635 (2016) (“[E]ven when the empirical scholars completely agree on the underlying facts, interpretation of the results can dramatically

Sixth, case outcomes are impacted by parties' factor-based calculation of a successful outcome. The Priest-Klein "selection hypothesis" predicts that, given various conditions, plaintiff win rates at trials should approach fifty percent because only the close cases survive settlement—or summary adjudication.⁶⁴ The hypothesis assumes parties have equal stakes in the litigation.⁶⁵ More recent studies cast the fifty-percent hypothesis in doubt, including those dealing specifically with intellectual property law.⁶⁶ Notably, "win rate" means the percentage of time one party (plaintiff or defendant) wins when a factor is decided in that party's favor, not the percentage of time that party wins when the factor is relevant. With these caveats in mind, this discussion turns to the theory underlying likelihood of confusion and the points of departure from conventional wisdom in practice.

B. *Blends, Triggers, and Polaroid Factors*

This Section opens by discussing the impact of blending technical trademarks and trade names in modern trademark law. It proceeds to introduce the likelihood of confusion factors before presenting a doctrinal and empirical analysis of intent, surveys, mark strength, and consumer sophistication, arguing that each, in turn, detracts from an accurate likelihood of confusion analysis. Finally, the Section closes by

differ. Empirical legal scholarship is still worth conducting, but the hope that it will resolve partisan debates in law is unrealistic.”).

64. See George L. Priest & Benjamin Klein, *The Selection of Disputes for Litigation*, 13 J. LEGAL STUD. 1, 4–6, 17–18 (1984). Priest and Klein's fifty percent has been modified when there are different stakes involved. For example, if the plaintiff has more to win than the defendant has to lose. See Yoon-Ho Alex Lee & Daniel M. Klerman, *Updating Priest and Klein* 2 (Univ. S. Cal. Ctr. L. & Soc. Sci., Research Paper No. 15-21, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2619856).

65. Priest & Klein, *supra* note 64, at 24–29.

66. See e.g., Theodore Eisenberg, *Testing the Selection Effect: A New Theoretical Framework with Empirical Tests*, 19 J. LEGAL STUD. 337, 338–39 (1990) (testing the fifty-percent hypothesis and rejecting it as a description of all civil litigation); see also Mark A. Lemley & Colleen V. Chien, *Are the U.S. Patent Priority Rules Really Necessary?*, 54 HASTINGS L.J. 1299 (2003) (arguing that the Priest-Klein hypothesis is not borne out by the data in patent cases); Jason Rantanen, *Why Priest-Klein Cannot Apply to Individual Issues in Patent Cases*, (Univ. Iowa Legal Stud., Research Paper No. 12-15, 2012), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2132810 (“At best, the Priest-Klein hypothesis only applies to the selection of disputes, not the selection of individual issues. Due to the presence of multiple issues in patent cases, there is axiomatically no basis for inferring that a patentee would expect a 50 percent chance of winning on each one.”).

explaining the impact of coherence-based reasoning on the likelihood of confusion factors.

1. *Blends and triggers*

Early trademark common law distinguished between trade names and technical trademarks.⁶⁷ Most trade name disputes involved rivals.⁶⁸ Unfair competition law governed these disputes and focused on directly competing uses diverting trade,⁶⁹ taking the form of passing off and reverse passing off business names.⁷⁰ In the twentieth century, courts blurred the distinction between technical trademarks and trade names, blending the most expansive aspects in favor of trademark owners.

In 1946, Congress “federalized” common law protection of trademarks used in interstate commerce with the Lanham Act.⁷¹ The Act codified this blended standard, requiring only that the unauthorized use be connected with goods or services.⁷² Trade names enjoyed the protection offered to technical trademarks as long as owners could show “secondary meaning.”⁷³ Cases interpreted this as an association by consumers with the source of the product that imbued trade names with an acquired distinctiveness.⁷⁴ The Act subsequently welded the two concepts, allowing all signs to acquire distinctiveness through secondary meaning.⁷⁵

67. See FRANK I. SCHECHTER, *THE HISTORICAL FOUNDATIONS OF THE LAW RELATING TO TRADE-MARKS* 161 (1925).

68. Edward S. Rogers, *The Lanham Act and the Social Function of Trademarks*, 14 *LAW & CONTEMP. PROB.* 173, 178–80 (1949).

69. See Mark P. McKenna, *The Normative Foundations of Trademark Law*, 82 *NOTRE DAME L. REV.* 1839, 1904 (2007) (noting “that courts only developed the likelihood of confusion factors after jettisoning the requirement of direct competition”).

70. See *CORPORATE COUNSEL’S GUIDE TO UNFAIR COMPETITION* § 3:31 (2021). “Passing off” occurs when defendants sell its goods with the plaintiff’s mark, whereas in “reverse passing off,” defendants sell plaintiff’s goods with the defendant’s trademark. *Id.*

71. *CAE, Inc. v. Clean Air Eng’g, Inc.*, 267 F.3d 660, 672 (7th Cir. 2001).

72. See Lanham Act, Pub. L. No. 79-489, § 32(1)(a), 60 Stat. 427, 437 (1946) (specifying that only goods and services fell under the purview of the Act).

73. See, e.g., Milton Handler & Charles Pickett, *Trade-Marks and Trade Names—An Analysis and Synthesis*, 30 *COLUM. L. REV.* 168, 168–69 (1930).

74. See *E. Columbia, Inc. v. Waldman*, 181 P.2d 865, 867 (Cal. 1947) (en banc).

75. Compare 15 U.S.C. § 1052(f) (2006) (allowing registration of a mark “which has become distinctive of the applicant’s goods in commerce”), with Trade-Mark Act of 1905, Pub. L. No. 58-84, § 5(b), 33 Stat. 724, 725–26 (providing that no mark that is distinguishable “from other goods of the same class shall be refused” trademark registration because of the nature of the mark).

As sellers expanded into adjacent product markets in the post-World War Two era, courts expanded the scope of protection to include complementary products and services.⁷⁶ Trademark scope could protect virtually anything that functions as an identifier of source—shapes, colors, smells, and sounds.⁷⁷ Congress included new types of protectable subject matter from technical trademarks to “anything . . . capable of carrying [source] meaning” as a potential trademark.⁷⁸ As a result, the likelihood of confusion standard became more complex. Although courts previously compared the marks, cases now require courts to consider a much broader range of information, including advertising slogans, product packaging, and product designs.

Figure 1: Trade Names, Technical Trademarks, and Modern Trademarks

	Trade Names	Technical Trademarks	Modern Trademarks
Distinctiveness	Sufficient if descriptive	Requires distinctiveness	Sufficient if descriptive
Intent	Intent required	Strict Liability	Intent optional
Harm	Actual harm required	Likelihood of harm sufficient	Likelihood of harm sufficient
Comparison	No	Yes	Optional
Injunction	Narrow	Broad	Broad

Contemporary empirical evidence from this Article’s dataset indicates that defendants win on the merits 26% of the time when the parties are rivals. Defendants win on the merits only 46% when they are not rivals, with rivals winning twice as often. These numbers show the impact of unfair competition in shaping modern trademark doctrine.

Congress subsequently amended the Act to remove the restriction on source confusion, allowing courts to consider other forms of

76. See PAMELA WALKER LAIRD, *ADVERTISING PROGRESS: AMERICAN BUSINESS AND THE RISE OF CONSUMER MARKETING* 31 (1998) (discussing post-war expansion of consumer products).

77. See Bone, *Notice Failure and Defenses*, *supra* note 11, at 1268.

78. *Qualitex Co. v. Jacobson Prods. Co.*, 514 U.S. 159, 162 (1995).

confusion in the infringement analysis.⁷⁹ Courts dutifully expanded the scope of confusion from purchasers to include non-purchasers (“post-sale confusion”) and allowed businesses to prohibit confusion over sponsorship or endorsement of goods and services.⁸⁰

Law and economics scholarship prompted this expansion, driven by a belief that stronger protection maximized wealth and, in turn, promoted economic efficiency.⁸¹ The resulting fusion infused unfair competition into trademark law and invited courts to find defendants’ marks infringing well before consumers purchased a product or service with the allegedly infringing mark, based on the idea that defendants misappropriated the plaintiff’s goodwill to appeal to consumers.⁸²

With new triggers, confusion can manifest itself in various ways. For instance, “forward confusion occurs when ‘the junior user attempts to trade on the senior[] user’s goodwill and reputation,’” misleading consumers to believe that the junior and senior user’s goods or services are related.⁸³ Similarly, reverse confusion occurs where consumers believe the junior user is the source of the senior user’s goods.⁸⁴ Whereas protection previously stopped at the shores of adjacent products, trademark law now allows even a pancake chain restaurant to attempt to prohibit an evangelical Christian organization from using

79. See S. REP. NO. 87-2107, at 4 (1962), *reprinted in* 1962 U.S.C.C.A.N. 2844, 2847.

80. Act of Oct. 9, 1962, Pub. L. No. 87-772, § 2, 76 Stat. 769, 769 (deleting the requirement that confusion be of purchasers as to the source of origin of such goods or services).

81. See *e.g.*, *W. T. Rogers Co. v. Keene*, 778 F.2d 334, 339 (7th Cir. 1985) (explaining that “competition is not impaired by giving each manufacturer a perpetual ‘monopoly’ of his identifying mark” if he has chosen a “distinctive” trademark where the available names are “for all practical purposes infinite”); see *Landes & Posner, supra* note 12, at 270–79 (advancing Chicago School economic theory within trademark law scope).

82. *Gibson Guitar Corp. v. Paul Reed Smith Guitars*, 423 F.3d 539, 549 (6th Cir. 2005).

83. *Fuel Clothing Co. v. Nike, Inc.*, 7 F. Supp. 3d 594, 610 (D.S.C. 2014) (quoting *Coryn Grp. II, LLC v. O.C. Seacrets, Inc.*, CIV WDQ-08-2764, 2010 WL 1375301, at *4 (D. Md. Mar. 30, 2010)).

84. *Kelly-Brown v. Winfrey*, 717 F.3d 295, 304–05 (2d Cir. 2013); *J.T. Colby & Co. v. Apple Inc.*, 586 F. App’x 8, 9 (2d Cir. 2014) (“The Lanham Act guards against this ‘reverse confusion’ to prevent ‘a larger, more powerful company [from] usurping the business identity of a smaller senior [trademark] user.’”).

a similar mark.⁸⁵ This development caused a jurisprudential disjuncture to occur.

While the statute had changed, earlier courts did not update the likelihood of confusion test, which had been designed to capture more than just source confusion.⁸⁶ As will be shown below, factors like consumer sophistication, the likelihood of expansion, and the marketing channels are of little assistance in evaluating whether a company's claim that it is the exclusive soda for sporting events in the minds of the consumers is true.⁸⁷ Worse, the multiple targets that the likelihood of confusion standard now addresses makes applying it even more unwieldy and unpredictable.⁸⁸

2. *The Polaroid factors*

The multifactor test for the likelihood of confusion attempts to provide analytical rigor to the complicated question of how consumers perceive different marks. Barton Beebe's 2006 empirical study revealed courts most frequently deployed the Second Circuit's test in *Polaroid Corp. v. Polaroid Electronics Corp.*⁸⁹ In that case, Judge Friendly articulated what became known as the eight *Polaroid* factors:

- (1) strength of the plaintiff's mark;
- (2) similarity of plaintiff's and defendant's marks;
- (3) competitive proximity of the products;
- (4) likelihood that plaintiff will "bridge the gap" and offer a product like a defendant's;
- (5) actual confusion between products;
- (6) good faith on the defendant's part;

85. First-Amended Complaint for Injunctive Relief for Federal Trademark Infringement and Dilution at 6, *IHOP IP, LLC v. Int'l House of Prayer*, No. CV10-6622 2010 WL 3775268 (C.D. Cal. Sept. 9, 2010).

86. *See, e.g.*, *King of the Mountain Sports, Inc. v. Chrysler Corp.*, 185 F.3d 1084, 1090 (10th Cir. 1999) (discussing the three-part test to determine similarity between marks).

87. Supreme Assembly, *Ord. of Rainbow for Girls v. J.H. Ray Jewelry Co.*, 676 F.2d 1079, 1082 (5th Cir. 1982) ("other association"); *Caterpillar Inc. v. Walt Disney Co.*, 287 F. Supp. 2d 913, 918 (C.D. Ill. 2003) ("otherwise affiliated").

88. James Gibson, *Risk Aversion and Rights Accretion in Intellectual Property Law*, 116 *YALE L.J.* 882, 908 (2007) ("The case law on sponsorship and approval, however, is so ambiguous as to make it almost impossible to know ex ante whether a given use will be infringing.").

89. Beebe, *supra* note 6, at 1593 ("This is especially true in the Second Circuit where the multifactor test is most often applied and where appellate panels have repeatedly emphasized that the multifactor analysis must be exhaustive and explicit."); *Polaroid Corp. v. Polarad Elecs. Corp.*, 287 F.2d 492, 495 (2d Cir. 1961).

- (7) quality of defendant's product; and
- (8) sophistication of the buyers.⁹⁰

Confusion is more likely when an accused product contains multiple indicia of similarity.⁹¹ For instance, house brands typically include house marks, product-specific marks, product packaging, and color or configuration.⁹² Conversely, consumers are less likely to be confused when defendants copy only a few elements.⁹³ However, no single factor in the likelihood of confusion inquiry is determinative. Conventional wisdom teaches that courts need to undertake "a highly fact-intensive inquiry both as to the assessment of the evidence concerning each factor and as to the overall synthesis of factors and the evidence."⁹⁴

Trademark litigation is inherently impressionistic, particularly because actual confusion is rare. Sometimes, each side claims a numerically equal number of factors in their favor, leaving courts to assign weights.⁹⁵ Courts caught up in the swirl sloppily pepper their judgments with different operative terms to describe the same thing, including affiliation,⁹⁶ endorsement,⁹⁷ connection,⁹⁸ and whether the use produced confusion "of any kind."⁹⁹ As the Fifth Circuit bluntly put it, "Congress adopted an open-ended concept of confusion. . . . Any kind of confusion will now support an action for trademark infringement."¹⁰⁰

Unfortunately, courts in subsequent cases as well as businesses and their legal advisors struggle to determine the appropriate strength of

90. *Polaroid*, 287 F.2d at 495.

91. See Beebe, *supra* note 6, at 1625 (noting that judges "emphasize similarities over differences," but finding that "the degree of similarity of [] marks does not appear to significantly affect the outcome of the test").

92. See, e.g., *Bristol-Myers Squibb Co. v. McNeil-P.P.C., Inc.*, 973 F.2d 1033, 1042 (2d Cir. 1992).

93. George Miaoulis & Nancy D'Amato, *Consumer Confusion & Trademark Infringement*, 42 J. MKTG. 48, 54 (1978) (finding, in the context of competing goods, that the "primary cue for [] association [between two brands] was not the name but the visual appearance").

94. *Select Comfort Corp. v. Baxter*, 996 F.3d 925, 933–34 (8th Cir. 2021) ("We have repeatedly emphasized that no one factor is controlling and different factors will carry more weight in different settings.")

95. *Equitable Nat'l Life Ins. Co. v. AXA Equitable Life Ins. Co.*, 434 F. Supp. 3d 1227, 1252 (D. Utah 2020) ("Ultimately, while each side can claim three factors weigh in its favor, they do not weigh equally.")

96. *Pebble Beach Co. v. Tour 18 I Ltd.*, 155 F.3d 526, 544 (5th Cir. 1998).

97. *Id.*

98. *Id.* at 543.

99. *Syntex Lab'ys, Inc. v. Norwich Pharmacal Co.*, 437 F.2d 566, 568 (2d Cir. 1971).

100. *Armstrong Cork Co. v. World Carpets, Inc.*, 597 F.2d 496, 500 n.5 (5th Cir. 1979).

each factor, either alone or relative to other factors.¹⁰¹ Judges themselves admit the distinctions they make are often done on an “intuitive basis” rather than through “logical analysis.”¹⁰² Reporting on his dataset of cases, Beebe observed that “scattered among the circuits are factors that are clearly obsolete, redundant, or irrelevant, or, in the hands of an experienced judge or litigator, notoriously pliable.”¹⁰³

Like an untended garden, the likelihood of confusion standard has grown wild, with different circuit courts spinning off anywhere between six and thirteen factors.¹⁰⁴ Some circuits favor factors others ignore, and courts have called nearly every factor or factor combination the most important.¹⁰⁵ The reason for this may be divergent conceptions of trademark policy, with some courts focusing on unfair competition while others are concentrating on consumer confusion.¹⁰⁶

This Article reveals for the first time in the trademark literature that the Second Circuit’s *Polaroid* factors no longer dominate modern trademark jurisprudence. Instead, the Ninth Circuit’s *Sleekcraft* factors have edged out the *Polaroid* factors as those most frequently applied as the Ninth Circuit now has the most trademark infringement cases.¹⁰⁷ As a result, the Second Circuit is now the second most dominant circuit.

101. Joseph P. Liu, *Two-Factor Fair Use?*, 31 COLUM. J.L. & ARTS 571, 579 (2008) (“Under a multi-factor balancing test, it is difficult to register the relative strength of the factors.”); Eric Goldman, *Online Word of Mouth and its Implications for Trademark Law*, in TRADEMARK LAW AND THEORY: A HANDBOOK OF CONTEMPORARY RESEARCH 404, 415–16, 424 (Graeme B. Dinwoodie & Mark D. Janis eds., 2008) (“Assessing consumer confusion about product source is an inherently inexact process.”).

102. *Union Carbide Corp. v. Ever-Ready Inc.*, 531 F.2d 366, 379 (7th Cir. 1976).

103. Beebe, *supra* note 6, at 1583–84.


104. *See infra*, Section I.B.

105. Beebe, *supra* note 6, at 1583.

106. Alejandro Mejías, *The Multifactor Test for Trademark Infringement from a European Perspective: A Path to Reform*, 54 IDEA 285, 314 (2014) (“[T]here is also divergence on how the factors are treated and employed.”); *see* Beebe, *supra* note 6, at 1596–97 (summarizing in chart form the different factors each circuit considers and reporting “substantial intercircuit variation in plaintiff multifactor test win rates.”).

107. *AMF Inc. v. Sleekcraft Boats*, 599 F.2d 341, 350–51 (9th Cir. 1979).


Figure 2: Circuit Variances over Time



This shift may be significant for litigation strategy, especially because the Second Circuit was the most defendant-friendly circuit where plaintiffs’ win rate was 31% and the defendant’s win rate was 48%. By comparison, both plaintiff and defendant win rates were 38% at the Ninth Circuit. In contrast to the Second and Ninth Circuits, the Federal Circuit was the most plaintiff-friendly, with an 83% plaintiff win rate. Defendants there fare comparatively poorly, winning a mere 8% of cases.¹⁰⁸ The figure below shows the distribution of cases and outcomes across circuits.

108. It is possible that defendants fare so poorly at the Federal Circuit because they are likely on appeal from the Trademark Trial and Appeal Board. However, the data is inconclusive on this point and invites further study.

Figure 3: Circuit by Outcome



As a doctrinal matter, the difference may be less material. As it turns out, the Ninth Circuit’s factors mirror those of the Second Circuit. The only difference between the two sets of factors is in linguistics as the Ninth Circuit considers marketing channels used to promote the products which is the same as the Second Circuit’s consideration of the competitive proximity of the products and services.¹⁰⁹ Similarly, this Article shows that the Second Circuit’s “quality of defendant’s product” factor can be subsumed into the competitive proximity factor.¹¹⁰ As an empirical matter, both factors rarely appear in case reports, with the Ninth Circuit’s “marketing channels” factor discussed in 13% of cases and the “quality” factor appearing in 11% of cases.¹¹¹

109. See e.g., *Las Vegas Sands Corp. v. Fan Yu Ming*, 360 F. Supp. 3d 1072, 1077 (D. Nev. 2019) (folding the two factors together).

110. See *infra* Section II.C.

111. Compare, for example, with mark similarity, which appeared in 85% of cases, competitive proximity in 73% of cases, and actual confusion in 74% of cases. See *infra* Appendix.

Figure 4: Comparing the Polaroid and Sleekcraft Factors

<i>Polaroid</i> Factors	<i>Sleekcraft</i> Factors
Strength of the plaintiff's mark	Strength of the plaintiff's mark
Similarity of plaintiff's and defendant's marks	Similarity of plaintiff's and defendant's marks
Competitive proximity of products or services	Competitive proximity of products or services
-	Marketing channels used to promote the products
Likelihood that plaintiff will "bridge the gap" and offer a product like a defendant's	Likelihood that either party will expand into new markets.
Actual confusion	Actual confusion
Defendant's good faith	Defendant's intent
Quality of defendant's product	-
Buyer sophistication	Type of goods and the carefulness of likely consumers

Substitution bias within each circuit's set of factors is particularly virulent when open-ended wording gives courts cover, as the Act does here.¹¹² Courts applying the Act took that opportunity and leaned into the likelihood of confusion factors like defendants' intent, survey evidence, and trademark strength, which were malleable and easy to wield to reach their desired outcomes.¹¹³ Savvy trademark attorneys also saw the opportunity to leverage more clever lawyering and focus less on the case's merits.¹¹⁴ Strikingly, Beebe's empirical study reported that intent and surveys were so heavily weighted that courts stampeded over other factors.¹¹⁵ Overreliance on these proxies results in a fundamentally flawed analysis. The next Section explains why.

112. See e.g., *SNA, Inc. v. Array*, 51 F. Supp. 2d 554, 562–63 (E.D. Pa. 1999), *aff'd sub nom. Silva v. Karlsen*, 259 F.3d 717 (3d Cir. 2001) (concluding that defendants' attempt to use metatags to "lure internet users to their site" was in bad faith).

113. See *infra* Section I.C.

114. See Beebe, *supra* note 6, at 1581 (suggesting this problem exists even with true source confusion cases because outcomes tend to be driven by the court's focus on intent).

115. *Id.* at 1607.

C. *Missing the Point on Consumer Confusion*

What constitutes “confusion” is highly subjective and difficult to evaluate. Proxies like intent, survey evidence, mark strength, and consumer sophistication fail to incorporate real-world purchasing conditions or are better considered within other factors. Trademark infringement is fundamentally flawed if the likelihood of confusion turns on these proxies.

1. *Intent*

Likelihood of confusion’s good faith or intent factor examines whether defendants sought to benefit from plaintiffs’ goodwill.¹¹⁶ All circuits but the Federal Circuit recognized this as a major factor in finding liability.¹¹⁷ “In analyzing whether a defendant has acted in bad faith, the question is whether the defendant attempted ‘to exploit the good will and reputation of a senior user by adopting the mark with the intent to sow confusion between the two companies’ products.’”¹¹⁸

Courts recognize that intentional copying may not indicate that the defendant attempted to capitalize on the plaintiff’s trademark or trade dress.¹¹⁹ However, there may be legitimate reasons to copy or imitate the primary features of another company’s product. These include functional features that have economic benefits without any secondary meaning.¹²⁰ In doing so, courts “want competitors to be inspired by—and to improve on—the findings of their predecessors.”¹²¹ Therefore, it is a “nefarious variety of passing off—the kind that confuses consumers and exploits a competitor’s established goodwill—that trademark law is prepared to prevent.”¹²²

Stating the distinction is easy in theory, hard in practice. Cases in the dataset reveal divergent views on when defendants cross the line. Some courts are prepared to exculpate defendants if they had no intent to confuse consumers.¹²³ Indeed, one court commended “upcycling,” or

116. *Sicilia Di R. Beibow & Co. v. Cox*, 732 F.2d 417, 431 (5th Cir. 1984) (discussing how the proper test focuses mainly on intent).

117. *See Beebe, supra* note 6, at 1589–90.

118. *Tiffany & Co. v. Costco Wholesale Corp.*, 971 F.3d 74, 88 (2d Cir. 2020).

119. *See Beebe, supra* note 6, at 1630.

120. *Fuddruckers, Inc. v. Doc’s B.R. Others, Inc.*, 826 F.2d 837, 844–45 (9th Cir. 1987).

121. *Vital Pharms., Inc. v. Monster Energy Co.*, No. 19-60809-CIV, 2021 WL 3371942, at *44 (S.D. Fla. Aug. 3, 2021).

122. *Id.*

123. *Id.* at *53 (“If a defendant intentionally copies an aspect of the plaintiff’s product, but not with intent to confuse consumers, then the defendant’s intent has

“restoring previously nonfunctional antique watch movements and parts,” as good faith.¹²⁴ The court used this reasoning despite the defendant’s intent to benefit from displaying the plaintiff’s mark, though intending to capitalize on its historical significance rather than its modern-day reputation.¹²⁵

Others courts stand ready to pin the defendant down on a lower negligence standard for failure to exercise due diligence.¹²⁶ Yet some will find against the defendant on an attempt standard, even without proof that actual confusion resulted from it,¹²⁷ and “some courts find evidence of bad faith even where they conclude the defendant did not choose its mark purposely to promote confusion.”¹²⁸ Unsurprisingly, these courts emphatically state that a defendant’s lack of intent is generally not relevant to the likelihood of consumer confusion.¹²⁹ Yet, surprisingly, the presence of intent may not be decisive either. For example, in one case, the court expressed that even when there is explicit evidence of bad faith, that factor alone should not determine the outcome of a case.¹³⁰ Instead, a defendant’s “[b]ad faith and intent to deceive are relevant to the extent that they add to the likelihood

little bearing on the ultimate question: whether the allegedly infringing product is likely to confuse consumers.”); *QuikTrip W., Inc. v. Weigel Stores, Inc.*, 984 F.3d 1031, 1036 (Fed. Cir. 2021) (“[T]he ‘only relevant intent is intent to confuse. There is a considerable difference between an intent to copy and an intent to deceive.’”).

124. *Hamilton Int’l Ltd. v. Vortic LLC*, 486 F. Supp. 3d 657, 667–68 (S.D.N.Y. 2020), *aff’d*, 13 F.4th 264, 268 (2d Cir. 2021) (“The Court credits this testimony, concluding that he did not intend to cause consumer confusion but rather sought to ‘preserve American history’ by salvaging and restoring the hearts of antique pocket watches.”).

125. *Id.* at 668; *see also* *Champion Spark Plug Co. v. Sanders*, 331 U.S. 125, 130 (1947) (noting that it can be “wholly permissible” that the “second-hand dealer gets some advantage from the trademark”).

126. *AWGI, LLC v. Atlas Trucking Co.*, 998 F.3d 258, 268 (6th Cir. 2021) (finding the intent factor irrelevant); *see also* *Ironhawk Techs., Inc. v. Dropbox, Inc.*, 994 F.3d 1107, 1124 (9th Cir. 2021), *amended by* 2 F.4th 1150 (9th Cir. 2021) (“This factor ‘favors the plaintiff where the alleged infringer adopted his mark with knowledge, actual or constructive, that it was another’s trademark.’”).

127. *ServPro Intell. Prop., Inc. v. Blanton*, 451 F. Supp. 3d 710, 727 (W.D. Ky. 2020) (“Courts have held that ‘[i]f a party chooses a mark with the intent of causing confusion, *that fact alone* may be sufficient to justify an inference of confusing similarity.’” (emphasis added) (quoting *Homeowners Grp., Inc. v. Home Mktg. Specialists, Inc.*, 931 F.2d 1100, 1111 (6th Cir. 1991))).

128. *Equitable Nat’l Life Ins. Co. v. AXA Equitable Life Ins. Co.*, 434 F. Supp. 3d 1227, 1248 (D. Utah 2020).

129. *ServPro Intell. Prop., Inc.*, 451 F. Supp. 3d, at 727; *Variety Stores, Inc. v. Wal-Mart Stores, Inc.*, 888 F.3d 651, 665 (4th Cir. 2021) (“Bad faith may be inferred from [Walmart’s] actual or constructive knowledge of [Variety’s] mark.”).

130. *Car-Freshner Corp. v. Am. Covers, LLC*, 980 F.3d 314, 333 (2d Cir. 2020).

that the accused infringer will achieve its objective of consumer confusion.”¹³¹

Beverly Pattishall suggested that factoring in intent makes outcomes more predictable.¹³² This inference makes it easier to determine the state of mind of one person, the defendant, than to forecast the perceptions of the consumer group.¹³³ Predictability is good, but the result may not be. Intent inherently focuses on the wrong goalpost. Merely because the defendant’s mental state is easier to discern than the perception of the consuming public does not make that factor more relevant to the inquiry. As Kelly Collins warned, “[t]his is dangerous because mere ‘copying’ is not always impermissible.”¹³⁴ The law encourages reusing generic or functional marks “as a part of our competitive economic system.”¹³⁵ For this reason, she argues that the relevant intent is the one to confuse and not merely to copy.¹³⁶

Another reason to abandon intent is that it muddies jurisprudential waters caused by further fusion of trade name and technical trademark jurisprudence. Courts typically require intent when dealing with non-inherently distinctive marks.¹³⁷ Courts have either presumed intent or dispensed with it for inherently distinctive marks.¹³⁸

Alejandro Mejías explained that intent is irrelevant because the focus “is not what the defendant intended to do, but whether his mark is likely to be confusingly similar for the relevant public.”¹³⁹ Very few courts acknowledge this much.¹⁴⁰ Judges may like intent because it makes their job easier, and the outcome feels more just. However,

131. *Id.*

132. See Beverly W. Pattishall, *The Impact of Intent in Trade Identity Cases*, 60 TRADEMARK REP. 575, 579–80 (1970).

133. *Id.* at 577.

134. Kelly Collins, Comment, *Intending to Confuse: Why Preponderance Is the Proper Burden of Proof for Intentional Trademark Infringements Under the Lanham Act*, 67 OKLA. L. REV. 73, 87 (2014).

135. *Id.*

136. *Id.* at 87–88 (“This would better serve the purposes of the Lanham Act and safeguard innocent conduct from triggering liability.”).

137. *But see* 4 MCCARTHY, *supra* note 2, § 23:106 (explaining that proof of intent is merely evidence relevant to whether confusion is likely).

138. See *Cashmere & Camel Hair Mfrs. Inst. v. Saks Fifth Ave.*, 284 F.3d 302, 317 (1st Cir. 2002) (explaining that where there is evidence of intentional deceit the presumption is clear); *Res. Devs., Inc. v. Statue of Liberty-Ellis Island Found., Inc.*, 926 F.2d 134, 140 (2d Cir. 1991).

139. Mejías, *supra* note 106, at 349.

140. See, e.g., *Virgin Enters. Ltd. v. Nawab*, 335 F.3d 141, 151 (2d Cir. 2003) (explaining that intent is not “of high relevance to the issue of likelihood of confusion” because “[i]t does not bear directly on whether consumers are likely to be confused”).

intent is irrelevant to technical trademark infringement. Recall from Section I.B that technical trademark infringement focuses on the consequences of the defendant's act and not on their intent.¹⁴¹ In contrast, trade name infringement focuses on defendants' desired outcomes, irrespective of consumer confusion.¹⁴²

Pinning the likelihood of confusion on free-riding becomes problematic because free-riding is ultimately a concept searching for meaning.¹⁴³ The Act does not require proof of intent. Trademark infringement is, after all, a strict liability offense.¹⁴⁴ As the Sixth Circuit opined, the better view is to consider intent only *after* other likelihood of confusion factors indicate liability.¹⁴⁵ Intent may go to aggravated remedies, but it should be irrelevant to the question of guilt. As Beebe put it, "if trademark law seeks to prevent commercial immorality, then it should do so explicitly. An injunction should issue, and damages be granted on that basis alone, and not on the basis of possibly distorted findings of fact as to the likelihood of consumer confusion."¹⁴⁶

Beebe found that despite the disconnect between the defendant's intent and consumer confusion, it stampedes the other factors.¹⁴⁷ The effect is powerful—a "nearly un-rebuttable presumption of a likelihood of confusion"¹⁴⁸ roughly 97% of the time,¹⁴⁹ making it "arguably the single most important confusion factor in use today."¹⁵⁰

141. Rogers, *supra* note 68, at 178 (explaining the origins of trademark law).

142. See, e.g., *Visible Sys. Corp. v. Unisys Corp.*, 551 F.3d 65, 73 (1st Cir. 2008) (explaining that sometimes there is a likelihood of confusion in industries regardless of intent).

143. See, e.g., *Ty Inc. v. Perryman*, 306 F.3d 509, 512 (7th Cir. 2002) (rejecting sponsorship dilution claim because "in that attenuated sense of free riding, almost everyone in business is free riding").

144. See *Taubman Co. v. Webfeats*, 319 F.3d 770, 775 (6th Cir. 2003) (recognizing that the Lanham Act is a "strict liability statute"); see also Rebecca Tushnet, *Running the Gamut From A to B: Federal Trademark and Federal False Advertising Law*, 159 U. PA. L. REV. 1305, 1310 (2011) (noting that federal courts have interpreted trademark as a strict liability offense); Robert G. Bone, *Enforcement Costs and Trademark Puzzles*, 90 VA. L. REV. 2099, 2109 (2004) (referring to trademark infringement as a form of strict liability).

145. See, e.g., *Taubman Co.*, 319 F.3d at 775 ("[T]he proper inquiry is not one of intent. In that sense, the Lanham Act is a strict liability statute. If consumers are confused by an infringing mark, the offender's motives are largely irrelevant.").

146. Beebe, *supra* note 6, at 1631.

147. *Id.* at 1621.

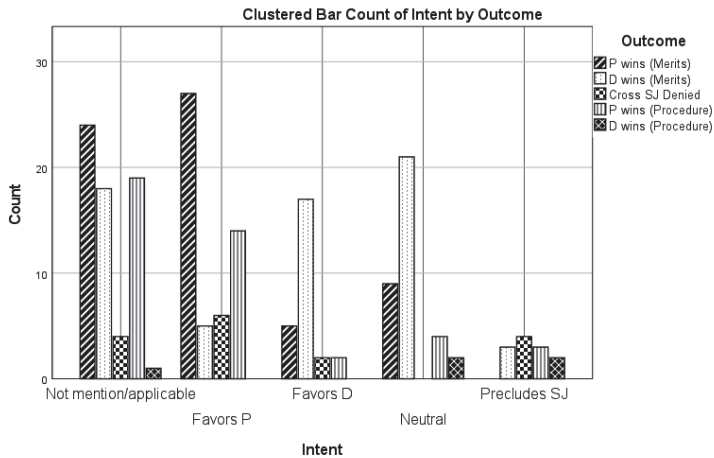
148. *Id.* at 1628.

149. *Id.* ("[The data] suggest that a finding of bad faith intent creates, if not in doctrine, then at least in practice, a nearly un-rebuttable presumption of a likelihood of confusion.").

150. Blake Tierney, *Missing the Mark: The Mispaced Reliance on Intent in Modern Trademark Law*, 19 TEX. INTEL. PROP. L.J. 229, 236 (2011).

This Article reports that intent appeared in two-thirds of the cases studied and was deemed neutral 19% of the time. In 27% of all cases, courts favored plaintiffs on the intent factor. When they did, plaintiffs won 52% of the time. In 14% of all cases, the courts favored defendants on the intent factor. When the court favored defendants, they won 65% of the time.

Figure 5: Intent by Outcome



Qualitatively, cases in the dataset show that intent bears a minimal impact on results.¹⁵¹ The reason is that “an intent to confuse customers is not required for a finding of trademark infringement.”¹⁵² Nonetheless, the intensely fact-specific nature of intent can trip up parties seeking speedy resolution of the dispute. As one court in the dataset put it, “[i]ssues of bad or good faith ‘are generally ill-suited for disposition on summary judgment.’”¹⁵³ The practice is longstanding,

151. *Equitable Nat’l Life Ins. Co. v. AXA Equitable Life Ins. Co.*, 434 F. Supp. 3d 1227, 1248 (D. Utah 2020) (“Although this factor weighs in AXA’s favor, its impact is minimal.”); *GoTo.com, Inc. v. Walt Disney Co.*, 202 F.3d 1199, 1208 (9th Cir. 2000) (emphasizing “the minimal importance of the intent factor”).

152. *GoTo.com*, 202 F.3d at 1208 (citation omitted).

153. *RVC Floor Decor, Ltd. v. Floor & Decor Outlets of Am., Inc.*, 527 F. Supp. 3d 305, 327 (E.D.N.Y. 2021); *see also Zamfir v. Casperlabs, LLC*, 528 F. Supp. 3d 1136, 1145 (S.D. Cal. 2021) (“These unresolved factual questions complicate the issue of Defendant’s intent in choosing the mark.”).

with courts preferring to leave it to juries to settle the matter.¹⁵⁴ Surprisingly, at least one court insisted on a jury trial even when the marks in question were identical due to the inherently subjective nature of the inquiry.¹⁵⁵

The dataset shows that 6% of cases expressly precluded summary judgment based on the intent factor. That figure may seem low, but it is considerably higher than any of the other factors: mark similarity (4%), buyer sophistication (3%), actual confusion (2%), mark strength (1%), competitive proximity (1%), “bridging the gap” (0%), and quality (0%).

Eliminating intent allows a more focused inquiry into the likelihood of confusion rather than the commercial immorality of defendants. As a practical matter, it frees parties from costly discovery and allows the court to grant summary judgment more frequently.¹⁵⁶ Judges can also dispose of cases more easily without trial, and it is less likely that defendants will be subject to vexatious suits based on the nebulous aspersions of intent.¹⁵⁷

While the “ordinary consumer” is central to the infringement analysis, it remains poorly theorized.¹⁵⁸ In patent cases, courts benefit from expert testimony.¹⁵⁹ Perhaps this is because the subject matter of patent disputes is by nature technologically challenging, defaulting those involved to accept, even expect, expert assistance. But trademark courts must investigate confusion without evidence that any consumers

154. See *EMI Catalogue P’ship v. Hill, Holliday, Connors, Cosmopolos Inc.*, 228 F.3d 56, 67–68 (2d Cir. 2000) (“Because the issue goes to defendants’ intent, it ‘is best left in the hands of the trier of fact.’”).

155. *Tiffany & Co. v. Costco Wholesale Corp.*, 971 F.3d 74, 88 (2d Cir. 2020) (“And as we have consistently observed, ‘subjective issues such as good faith are singularly inappropriate for determination on summary judgment.’”).

156. 10B CHARLES ALAN WRIGHT ET AL., *FEDERAL PRACTICE AND PROCEDURE (CIVIL)* § 2730 n.3 (3d ed. 2015) (“Questions of intent, which involve intangible factors including witness credibility, are matters for the consideration of the fact finder after a full trial and are not for resolution by summary judgment.”).

157. Thomas L. Casagrande, *A Verdict for Your Thoughts? Why an Accused Trademark Infringer’s Intent Has No Place in Likelihood of Confusion Analysis*, 101 TRADEMARK REP. 1447, 1455 (2011) (proposing an elimination of intent as a factor to be considered in determining trademark infringement).

158. See e.g., Thomas R. Lee et al., *Trademarks, Consumer Psychology, and the Sophisticated Consumer*, 57 EMORY L.J. 575, 575 (2008) (“[N]either courts nor commentators have made any serious attempt to develop a framework for understanding the conditions that may affect the attention that can be expected to be given to a particular purchase.”).

159. Douglas G. Smith, *The Increasing Use of Challenges to Expert Evidence Under Daubert and Rule 702 in Patent Litigation*, 22 J. INTELL. PROP. L. 345, 354 (2015).

were confused, imagining consumers' likely experience as filtered through the parties' competing interests. This notional consumer is "neither savant nor dolt."¹⁶⁰ One who "lacks special competency with reference to the matter at hand but has and exercises a normal measure of the layman's common sense and judgment."¹⁶¹ Instead, courts rely on surveys, mark strength, and consumer sophistication to determine the likelihood of confusion. But, like intent, none of these factors provide a good proxy. The Sections below explain why.

2. Surveys

Surveys attempt to measure whether consumers believe that the plaintiff's mark is the source of the alleged infringer's product or whether it sponsors or approves it.¹⁶² Plaintiffs may provide survey evidence that an appreciable number of relevant consumers are likely to be confused.¹⁶³ According to a case in the dataset, survey evidence is not a prerequisite for establishing public recognition, but "it is the most persuasive evidence of it."¹⁶⁴

Surveys present respondents with defendants' marks and measure consumers' reactions in the context that consumers encounter the mark in question.¹⁶⁵ Proof of marketing supports broad public recognition.¹⁶⁶ They typically involve control groups to show causality between the defendants' mark and consumer confusion.¹⁶⁷

In theory, a survey needs to pass muster under the Federal Rules of Evidence, which requires considering the "validity of the techniques employed."¹⁶⁸ Courts can bar significantly flawed surveys as evidence

160. Kraft Foods Grp. Brands LLC v. Cracker Barrel Old Country Store, Inc., 735 F.3d 735, 743 (7th Cir. 2013).

161. United States v. 88 Cases, More or Less, Containing Bireley's Orange Beverage, 187 F.2d 967, 971 (3d Cir. 1951).

162. 3 ANNE GILSON LALONDE, GILSON ON TRADEMARKS § 8.03 (2021).

163. See 6 MCCARTHY, *supra* note 2, § 32:158.

164. Kibler v. Hall, 843 F.3d 1068, 1074 (6th Cir. 2016).

165. Shari Seidman Diamond & David J. Franklyn, *Trademark Surveys: An Undulating Path*, 92 TEX. L. REV. 2029, 2037 (2014).

166. Therma-Scan, Inc. v. Thermoscan, Inc., 295 F.3d 623, 632 (6th Cir. 2002).

167. See, e.g., Bracco Diagnostics, Inc. v. Amersham Health, Inc., 627 F. Supp. 2d 384, 448 (D.N.J. 2009) (criticizing a survey's design for failure to use "a control mechanism").

168. Shari Seidman Diamond, *Reference Guide on Survey Research*, in REFERENCE MANUAL ON SCI. EVIDENCE 359, 364 (3d ed. 2011).

when they are more prejudicial than probative¹⁶⁹ or deemed unreliable.¹⁷⁰

The problem is that commentators and courts alike acknowledge that surveys are often unreliable and expensive, costing hundreds of thousands of dollars.¹⁷¹ Courts routinely attack the representativeness of the survey from a parade of cherry-picked witnesses and extrapolate a standard of what consumers generally believe.¹⁷² The inexact science of assessing trademark strength causes judges to rely upon or reject surveys based on whether the results agree with their subjective impressions.¹⁷³ As a result, judicial unease with surveys sometimes bubbles to the surface, with Judge Richard Posner remarking that “no doubt there are other tricks of the survey researcher’s black arts that we have missed.”¹⁷⁴

Constructing a robust survey is dauntingly hard. Surveys need to employ a control¹⁷⁵ and calculate noise.¹⁷⁶ As an indication of the

169. *Citizens Fin. Group, Inc. v. Citizens Nat’l Bank*, 383 F.3d 110, 120 (3d Cir. 2004).

170. *Id.*

171. *Nabisco, Inc. v. PF Brands, Inc.*, 191 F.3d 208, 224 (2d Cir. 1999) (“[C]onsumer surveys . . . are expensive, time-consuming and not immune to manipulation.”); *Johnson v. Revenue Mgmt. Corp.*, 169 F.3d 1057, 1063 (7th Cir. 1999) (“[S]urvey evidence in trademark and trade dress cases can be very costly.”); see also Robert H. Thornburg, *Trademark Survey Evidence: Review of Current Trends in the Ninth Circuit*, 21 SANTA CLARA COMPUT. & HIGH TECH. L.J. 715, 717 (2005) (“[T]he most basic of surveys cost[s] in the hundreds of thousands of dollars.”).

172. *But cf. Citizens Fin. Grp.*, 383 F.3d at 122 (“In general, ‘actual confusion’ evidence collected by employees of a party in a trademark action must be viewed with skepticism because it tends to be biased or self-serving.”).

173. 6 MCCARTHY, *supra* note 2, § 32:196 (“Since an estimation of the probable mental reactions and associations of the buying public is not a science, there is always the temptation to decide on the basis of a ‘hunch.’ That is, the trier of fact (or any human being) would rather extrapolate from his or her own subjective impressions than extrapolate from some hard evidence of other persons’ subjective impressions—especially if the two do not agree.”); see also Peter Weiss, *The Use of Survey Evidence in Trademark Litigation: Science, Art or Confidence Game?*, 80 TRADEMARK REP. 71, 83 (1990) (“[A] reading of the many cases in which either great weight or little weight was given to survey evidence will, I feel reasonably certain, lead most objective analysts to the conclusion that, while some surveys went down because they were indeed ‘seriously flawed,’ many others either stayed up or went down depending on the result which the judges wanted to reach.”).

174. *Indianapolis Colts, Inc. v. Metro. Balt. Football Club Ltd. P’ship*, 34 F.3d 410, 416 (7th Cir. 1994).

175. 6 MCCARTHY, *supra* note 2, § 32:187 (“Courts have held that a survey that fails to use a control may be given less weight or even excluded from evidence altogether.”).

176. DANIEL KAHNEMAN, OLIVER SIBONY & CASS SUNSTEIN, *NOISE: A FLAW IN HUMAN JUDGMENT* 488 (2021) (“Noise is variability in judgments that should be identical.”).

treacherousness of this task, one court dismissed an expert witness who authored a book on the very subject of conducting trademark surveys for producing a “useless” survey.¹⁷⁷ As a matter of justice between the parties, the staggering costs of surveys put defendants at a disadvantage. Robert Bone explained that “[p]roving a high [likelihood of confusion] puts a premium on surveys and expert testimony and is likely to require extensive discovery, all of which will increase direct litigation costs and strengthen a trademark owner’s ability to leverage cease-and-desist threats in frivolous and weak cases.”¹⁷⁸

Qualitatively, cases in the dataset warn that surveys only represent circumstantial evidence of actual confusion, providing an experimental environment, not real consumers making mistaken purchases.¹⁷⁹ As one court put it, “[a]necdotal evidence can be more direct evidence of actual confusion and so is ‘both relevant and probative.’”¹⁸⁰ Another court “noted a trend away from according great weight to survey evidence,” and afforded the survey no weight.¹⁸¹ Unlike actual confusion, that court explained that “survey evidence is circumstantial, not direct, evidence of the likelihood of confusion. Surveys do not measure the degree of actual confusion by real consumers making mistaken purchases.”¹⁸² Accordingly, the court faulted the survey for “depart[ing] from real-market conditions in a way that was both biased and misleading.”¹⁸³

177. *Vital Pharms., Inc. v. Monster Energy Co.*, No. 19-60809-CIV, 2021 WL 3371942, at *65 (S.D. Fla. Aug. 3, 2021).

178. See Bone, *Notice Failure and Defenses*, *supra* note 11, at 1269 n.110.

179. *Vital Pharms., Inc.*, 2021 WL 3371942, at *55 (“[S]urvey evidence is not direct evidence of customer confusion in the real marketplace.”).

180. *Juul Labs, Inc. v. 4X PODS*, 509 F. Supp. 3d 52, 66 (D.N.J. 2020).


181. *Vital Pharms., Inc.*, 2021 WL 3371942, at *55.

182. *Id.* at *64.

183. *Id.* at *65 (“[E]vidence at trial confirmed the obvious: that the artificial coolers Mr. Berger showed his survey participants looked nothing like the coolers consumers would encounter in real stores.”); see, e.g., *Citizens Banking Corp. v. Citizens Fin. Grp., Inc.*, 320 F. App’x 341, 348 n.4 (6th Cir. 2009) (minimizing the weight of a confusion survey because it “failed to mimic the purchase conditions”); *Coherent, Inc. v. Coherent Techs., Inc.*, 935 F.2d 1122, 1126 (10th Cir. 1991) (affirming the district court’s finding that the “survey did not show actual confusion because it failed to simulate decisions in the marketplace”); *Am. Footwear Corp. v. Gen. Footwear Co.*, 609 F.2d 655, 661 (2d Cir. 1979) (finding that “the critical defect in this survey was the failure to conduct it under actual marketing conditions”—and so the “district court’s rejection of this survey evidence was not clearly erroneous”); 4 MCCARTHY, *supra* note 2, § 23:2.50 (stating that a survey is only evidence of confusion if “the survey mirrors the real world setting which can create an instance of actual confusion”).

Scholars also warn against placing a premium on surveys. According to Beebe, “the conventional view of the utility of survey evidence may be incorrect”: only 20% of the cases he reviewed addressed survey evidence, 10% credited survey evidence, and 7% ruled in favor of the outcome that the credited survey evidence favored.¹⁸⁴ This dataset shows a near-identical result fifteen years later. Of the 20% of cases that addressed survey evidence, 12% credited survey evidence, and 6% ruled in favor of the outcome that the credited survey evidence favored.

Figure 6: Outcome by Survey Evidence



As with intent, there is a certain circular irony to the whole exercise regarding surveys. Courts rely on surveys only to support conclusions that judges reach using other factors. The analysis also works backward—faced with survey evidence showing a likelihood of confusion, judges may regard the marks as more similar than they might have appeared in the absence of the survey.¹⁸⁵ As Peter Weiss remarked, “[o]ne might sum it all up by saying that the function of surveys in trademark litigation is to plumb the minds of the public to make up the minds of the judges.”¹⁸⁶ Dispensing with surveys and

184. Beebe, *supra* note 6, at 1641. A more recent study that expanded Beebe’s dataset found that only about 17% of cases addressed survey evidence. See Robert C. Bird & Joel H. Steckel, *The Role of Consumer Surveys in Trademark Infringement: Empirical Evidence from the Federal Courts*, 14 U. PA. J. BUS. L. 1013, 1035 (2012).

185. Diamond & Franklyn, *supra* note 165, at 2043.

186. Weiss, *supra* note 173, at 86.

relying on the court's judgment would not only be cheaper and simpler, but it would also be the intellectually honest thing to do.

Surveys sometimes overlap with trademark strength since parties may use the former to measure the potency of a mark's goodwill and its worthiness of protection.¹⁸⁷ Known as the *Abercrombie* spectrum, generic and descriptive marks are not distinctive, suggestive marks are marginally distinctive, while arbitrary or fanciful marks are inherently distinctive.¹⁸⁸ Trademark strength is usually the first factor courts consider.¹⁸⁹

3. *Mark strength*

A mark's distinctiveness is its uniqueness in denoting a product. Marks may be fanciful, arbitrary, suggestive, descriptive, or generic from most to least distinctive.¹⁹⁰ Generic terms are unprotectable and descriptive ones are protectable only when buyers view them as distinctive of a unique source.¹⁹¹ Evaluating the strength of a mark requires the fact finder to evaluate several factors: its degree of inherent distinctiveness, its "conceptual strength," its distinctiveness in the marketplace, and its "commercial strength."¹⁹² Unlike conceptual strength, commercial strength considers advertising expenditures, consumer studies linking the mark to a source, sales success, unsolicited media coverage of the product, attempts to plagiarize the mark, and the length and exclusivity of the mark's use.¹⁹³

The dataset reveals that mark strength comes up in 70% of the cases. In 47% of all cases, the courts favored plaintiffs on the mark strength factor. When they did, plaintiffs won 46% of the time. In 15% of all

187. Beebe, *supra* note 6, at 1646 ("In trademark law, the question is always of consumer perception in the marketplace rather than judicial perception in the courtroom.").

188. *Abercrombie & Fitch Co. v. Hunting World, Inc.*, 537 F.2d 4, 9 (2d Cir. 1976).

189. *See, e.g., Welding Servs., Inc. v. Forman*, 509 F.3d 1351, 1361 (11th Cir. 2007) ("The stronger or more distinctive a trademark or service mark, the greater the likelihood of confusion . . ."); Barton Beebe & C. Scott Hemphill, *The Scope of Strong Marks: Should Trademark Law Protect the Strong More Than the Weak?*, 92 N.Y.U. L. REV. 1339, 1349 n.40 (2017) ("Strength is the first factor in the Second, Fourth, Fifth, Sixth, Eighth, Ninth, and Eleventh Circuits, the second factor in the Third Circuit, and the last factor in the First and Tenth Circuits."). Courts consider design marks under the *Seabrook* factors. *See Seabrook Foods, Inc. v. Bar-Well Foods, Ltd.*, 568 F.2d 1342 (C.C.P.A. 1977).

190. *U.S. Pat. & Trademark Off. v. Booking.com B. V.*, 140 S. Ct. 2298, 2302 (2020).


191. *Id.* at 2303.

192. Ouellette, *supra* note 40, at 353.

193. *Variety Stores, Inc. v. Walmart Inc.*, 852 F. App'x 711, 719 (4th Cir. 2021).

cases, the courts favored defendants on the mark strength factor. When the court favored defendants, they won 71% of the time.

Figure 7: Mark Strength by Outcome



In an empirical study on mark strength, Lisa Ouellette observed that “courts often have difficulty applying these tests.”¹⁹⁴ According to her, [t]he complex doctrine that has evolved around trademark strength and the likelihood of confusion appears to be a (largely unsuccessful) attempt to provide some analytical rigor to the essential questions of how strongly a mark identifies goods or services and how well it distinguishes those products from others in the marketplace.¹⁹⁵

Determining the bounds of an owner’s trademark requires more than just looking at the mark; it requires assessing what protection the trademark owner should be entitled to for that mark.¹⁹⁶ Distinctive marks are memorable as source indicators and possess greater conceptual strength to consumers.¹⁹⁷ Courts equate distinctiveness with a greater breadth of protection, are more willing to find confusing

194. Ouellette, *supra* note 40, at 353.

195. *Id.* at 360.

196. See *Virgin Enters. Ltd. v. Nawab*, 335 F.3d 141, 148 (2d Cir. 2003).

197. See *id.* (noting that consumers are more likely to attribute two products with more unique names to the same source versus two products with more generic names).

similarities,¹⁹⁸ and usually find that the strongest marks merit the widest range of protection.¹⁹⁹

Like the “black arts” of surveys, empirical studies confirm that courts judge mark strength intuitively.²⁰⁰ For instance, Beebe reported how courts failed to categorize the plaintiff’s mark in a specific category of distinctiveness in half of the cases he studied.²⁰¹ He observed that “considerations such as the comparative quality of the parties’ goods or the inherent distinctiveness of the plaintiff’s mark rarely aid in this inquiry.”²⁰² Others have variously criticized trademark strength as “needlessly open-ended”²⁰³ and “inconsistent.”²⁰⁴ One court acknowledged distinctiveness “is far from an exact science and that the differences between the classes, which is not always readily apparent . . . makes placing a mark in its proper context . . . tricky business at best.”²⁰⁵

As with survey evidence, Thomas McCarthy notes, that

a cynic would say that . . . when the court wants to find no infringement, it says that the average buyer is cautious and careful But if the judge thinks there is infringement, the judge sets the standard lower and says the average buyer is gullible and not so discerning.²⁰⁶

198. See, e.g., *First Sav. Bank, F.S.B. v. First Bank Sys., Inc.*, 101 F.3d 645, 655 (10th Cir. 1996) (“When the primary term is weakly protected to begin with, minor alterations may effectively negate any confusing similarity between the two marks.”).

199. See, e.g., *Ford Motor Co. v. Money Makers Auto. Surplus, Inc.*, No. 03CV493, 2005 WL 2464715, at *1, *3 (D. Neb. Sept. 14, 2005) (finding that the various Ford Motor Company marks at issue “are among the most famous marks in the world” and are “therefore entitled to the widest scope of protection”).

200. See, e.g., Thomas R. Lee, Eric D. DeRosia & Glenn L. Christensen, *Sophistication, Bridging the Gap, and the Likelihood of Confusion: An Empirical and Theoretical Analysis*, 98 TRADEMARK REP. 913, 913 (2008) (analyzing how courts rely on “precedent built on ‘personal intuition and subjective, internalized, stereotypes.’”); see also Beebe, *supra* note 6, at 1581 (describing the variation among circuits in their application of multifactor tests for likelihood of confusion).

201. Beebe, *supra* note 6, at 1633–35 (stating that some use of the spectrum was made in only 193 out of 331 cases and that the mark was placed in a specific category in only 164 cases).

202. *Id.* at 1645.

203. Timothy Denny Greene & Jeff Wilkerson, *Understanding Trademark Strength*, 16 STAN. TECH. L. REV. 535, 582 (2013).

204. Beebe, *supra* note 6, at 1633.

205. *Banff, Ltd. v. Federated Dep’t Stores, Inc.*, 841 F.2d 486, 489 (2d Cir. 1988).

206. 4 MCCARTHY, *supra* note 2, § 23:92; see also Ann Bartow, *Likelihood of Confusion*, 41 SAN DIEGO L. REV. 721, 747 (2004) (noting that judges give meaning to terms on a “case-by-case” basis).

The courts themselves regard likelihood of confusion merely “as a heuristic device to assist in determining whether confusion exists.”²⁰⁷ There is no requirement for the likelihood of confusion to consider survey evidence or mark strength. Eliminating both would simplify the likelihood of confusion and make it less prone to error.

4. *Consumer sophistication*

Consumer sophistication provides context to the consumer information available and the ability of consumers to discern between the marks.²⁰⁸ Courts analyze the degree of care reasonably expected of potential customers from the perspective of “the ordinary purchaser, buying under the normally prevalent conditions of the market and giving the attention such purchasers usually give in buying that class of goods.”²⁰⁹ More expensive products or services mean consumers take more time and effort when making decisions, and therefore, the likelihood of confusion decreases.²¹⁰ However, the defendant’s distribution methods may affect consumers’ degree of care, even when an individual product is not expensive.²¹¹

Scholars criticized the artificiality of consumer sophistication, likening it to expecting judges to perform a “Vulcan mind-meld” with consumers in the marketplace.²¹² Courts may easily project their normative view of how careful a consumer should be or their view of a

207. *Sullivan v. CBS Corp.*, 385 F.3d 772, 778 (7th Cir. 2004).

208. Andrew Martineau, *Imagined Consumers: How Judicial Assumptions About the American Consumer Impact Trademark Rights, for Better and for Worse*, 22 DEPAUL J. ART, TECH. & INTELL. PROP. L. 337, 352 (2012) (“This would seem to be a crucial part of the test, given that the standard for infringement is whether consumers are likely to be confused.”).

209. *Gen. Mills, Inc. v. Kellogg Co.*, 824 F.2d 622, 627 (8th Cir. 1987).

210. *Kibler v. Hall*, 843 F.3d 1068, 1080 (6th Cir. 2016) (observing that when consumers exercise caution in purchasing items, they are less likely to confuse their origins, such as “when consumers have expertise in the items and when the items are particularly expensive”).


211. *See, e.g., ZW USA, Inc. v. PWD Sys., LLC*, 889 F.3d 441, 447–48 (8th Cir. 2018) (finding fact that parties sold their respective low-cost products on different websites under different trade names strongly cut against a likelihood of confusion).

212. *See* William E. Gallagher & Ronald C. Goodstein, *Inference Versus Speculation in Trademark Infringement Litigation: Abandoning the Fiction of the Vulcan Mind Meld*, 94 TRADEMARK REP. 1229, 1230 (2004) (criticizing the lack of empirical evidence required to validate an inference of likelihood of confusion, such as no requirement for consumer surveys or evidence showing actual confusion).

defendant's conduct.²¹³ But, like intent, surveys, and mark strength, consumer sophistication suffers from inherent capriciousness.

The dataset reveals that consumer sophistication comes up in 46% of the cases, among the lowest of all the *Polaroid* factors. In 18% of all cases, the courts favored plaintiffs on the consumer sophistication factor. When they did, plaintiffs won 49% of the time. In 14% of all cases, the courts favored defendants on the consumer sophistication factor. When this factor favored defendants, they won 63% of the time.

Figure 8: Buyer Sophistication by Outcome



Three irrelevant factors are plenty, but there is one final culprit. That is, the sheer multitude of factors courts must consider. The total number of factors makes the likelihood of confusion analysis difficult to deploy, bogging down courts and encouraging selective application. Instead, judges and juries rely on coherence-based reasoning to make sense of their findings to cope with the sheer number of factors.

D. Coherence-Based Reasoning

Over the past century, trademark law ossified the likelihood of confusion standard from pragmatic judge-made rules of thumb into a rigid and formalistic standard. The *Restatement (First) of Torts* merely

213. *August Storck K.G. v. Nabisco, Inc.*, 59 F.3d 616, 618 (7th Cir. 1995) (“Many consumers are ignorant or inattentive, so some are bound to misunderstand no matter how careful a producer is.”).

mentioned “the following factors are important,” and the early cases applied the factors loosely.²¹⁴ However, appeals courts chastised lower courts for failing to address each factor, with orders to reverse and remand.²¹⁵ We can deduce this formalism ended up burdening courts with an unwieldy craft, forcing judges to pay lip service to all the factors while systemically relying on only a few. At the same time, their opinions recite disclaimers that the likelihood of confusion factors act only as a guide and that no single factor is dispositive.

Studies show that experts do not integrate multifactor test (“MFT”) factors well.²¹⁶ Even using stringent tests to aid in decision-making can lead to consistent and predictable mistakes.²¹⁷ It may occur early in the decision-making process, and a single attribute can trigger coherence-based reasoning.²¹⁸

Trademark law expects courts to decipher between six and thirteen likelihood of confusion factors, which often point in opposite directions, yet still reach a coherent conclusion in every case.²¹⁹ Worse, the likelihood of confusion factors in each circuit are not exhaustive, with courts occasionally considering other factors such as geographical proximity.²²⁰

Courts are divided on whether “it is incumbent upon the district judge to engage in a deliberate review of each factor.”²²¹ Some emphatically state

214. RESTATEMENT (FIRST) OF TORTS § 729 (AM. L. INST. 1939).

215. Barton Beebe, *An Empirical Study of the Multifactor Tests for Trademark Infringement*, 94 CALIF. L. REV. 1581, 1593 (2006).

216. See, e.g., Robyn M. Dawes, *The Robust Beauty of Improper Linear Models in Decision Making*, 34 AM. PSYCH. 571, 573 (1979) (positing that experts in a field are better at selecting and coding information than integrating it).

217. See Chris Guthrie et al., *Inside the Judicial Mind*, 86 CORNELL L. REV. 777, 779–80 (2001).

218. See Dan Simon, Daniel C. Krawczyk, & Keith J. Holyoak, *Construction of Preferences by Constraint Satisfaction*, 15 PSYCH. SCI. 331, 331 (2004) (suggesting that a single variable can initiate spreading coherence).

219. See *Tana v. Dantanna's*, 611 F.3d 767, 775 n.7 (11th Cir. 2010) (explaining that this test “presupposes that various factors will point in opposing directions”; it is the job of the court to determine the relative importance of the evidence probative of each factor in an effort to decide whether, “in light of the evidence as a whole, there is sufficient proof of a likelihood of confusion to warrant a trial of the issue”).

220. See *id.* at 781 (holding that “[t]he district court did not err in considering the geographic proximity of use as an eighth factor demonstrating the unlikelihood of confusion”); see also *J-B Weld Co. v. Gorilla Glue Co.*, 978 F.3d 778, 789 (11th Cir. 2020) (“While all seven factors must be considered, they are not necessarily exhaustive if other evidence is probative of a likelihood of confusion.”).

221. Compare *Arrow Fastener Co. v. Stanley Works*, 59 F.3d 384, 400 (2d Cir. 1995) (discussing each factor), with *Bumble Bee Seafoods LLC v. UFS Indus., Inc.*, No. 04

that “the factors are not truly independent—depending on the context, a strong showing as to one factor may serve to make a different factor more or less important.”²²² Yet others rule only on a few key factors, allowing them to resolve the dispute without needing a trial.²²³

Without meaningful guidance, courts weigh those factors impressionistically. Beebe’s study confirms that judges in the likelihood of confusion cases employ “‘fast and frugal’ heuristics to short-circuit the multifactor analysis.”²²⁴ Coherence-based reasoning operates bidirectionally to fit together how a judge decides the factors,²²⁵ both preceding the decision and in forming its basis.²²⁶ In other words, fact-finders assessing a likelihood of confusion test will look at the evidence as non-independently relative to the final decision.²²⁷ Consequently, the resulting decision is biased because, as Dan Simon explains, “the hard case morphs into an easy one” in the mind of the fact-finder.²²⁸

Formulating optimal legal rules requires judges to balance factors while taking account of “false positive” errors (i.e., prohibiting beneficial conduct) versus “false negative” errors (i.e., permitting harmful conduct).²²⁹ This task requires judges to access information on the frequency and impact of the error, the likelihood of deterrence,

Civ. 2105, 2004 WL 1637017, at *7 (S.D.N.Y. July 20, 2004) (discussing only the relevant factors).

222. *Select Comfort Corp. v. Baxter*, 996 F.3d 925, 934 (8th Cir. 2021).

223. *See, e.g., Hamilton Int’l Ltd. v. Vortic LLC*, 486 F. Supp. 3d 657, 666 (S.D.N.Y. 2020) (“As the Court noted in its summary judgment opinion, a number of the *Polaroid* factors are not helpful to this case.”), *aff’d*, 13 F.4th 264 (2d Cir. 2021).

224. Beebe, *supra* note 6, at 1581; Tierney, *supra* note 150, at 235–36 (“[M]uch of the time spent going through the list of factors in any given case is in reality just an attempt to justify a predetermined conclusion about the likelihood of confusion . . .”).

225. Dan Simon, *A Third View of the Black Box: Cognitive Coherence in Legal Decision Making*, 71 U. CHI. L. REV. 511, 515–16 (2004) [hereinafter Simon, *Third View of the Black Box*].

226. *See, e.g., Dan Simon et al., The Redux of Cognitive Consistency Theories: Evidence Judgments by Constraint Satisfaction*, 86 J. PERSONALITY & SOC. PSYCH. 814, 816 (2004).

227. Dan Simon, *The Limited Diagnosticity of Criminal Trials*, 64 VAND. L. REV. 143, 195 (2011).

228. Simon, *Third View of the Black Box*, *supra* note 225, at 517 (describing studies where coherence-based reasoning caused subjects who found for the defendant and those who found for the plaintiff to be more confident the evidence supported their view after they had issued their verdict).

229. Andrew I. Gavil & Steven C. Salop, *Probability, Presumptions and Evidentiary Burdens in Antitrust Analysis: Revitalizing the Rule of Reason for Exclusionary Conduct*, 168 U. PA. L. REV. 2107, 2119 (2020).

and the cost to the administrative process.²³⁰ On occasion, courts themselves express frustration with the likelihood of confusion tests, acknowledging that “[a]lthough our test for a likelihood of confusion is well-developed, some uncertainty remains as to *when* confusion must exist in order to support a trademark infringement claim.”²³¹

The takeaway is that an overload of factors demands too much from judges and forces them to stampede over those they deem less significant. In the absence of direct evidence of confusion, courts must ascertain it through a host of proxy factors.²³² Under these trying circumstances, Beebe empirically observed intent and actual confusion playing an outsized role in coloring how courts treated the other likelihood of confusion factors, confirming their perniciousness.²³³

As Michael Grynberg noted, “[e]ven if judges do no more than applying heuristics of questionable quality to the disposition of trademark claims, channeling the process through a consistent framework aids litigants in identifying and accommodating the factors that guide fact finding.”²³⁴ The question then is, how many factors should we retain?

Only a few, argued Beebe, pointing out that judges in the likelihood of confusion cases find only a few factors probative anyway.²³⁵ Indeed, cases in the dataset recognize that courts can short circuit the process and focus on just a few factors.²³⁶ Beebe recommended three or four

230. See *id.* at 2119–20 (asserting that formulating an optimal legal standard involves considering error costs, deterrence, and administrative costs).

231. See, e.g., *Select Comfort Corp. v. Baxter*, 996 F.3d 925, 934 (8th Cir. 2021).

232. Laura A. Heymann, *The Reasonable Person in Trademark Law*, 52 ST. LOUIS U. L.J. 781, 783 (2008).

233. Beebe, *supra* note 6, at 1620–21 (“Intent and, to a lesser degree, actual confusion appear to exert such a coherence-shifting influence when they favor a likelihood of confusion. Indeed, in the forty-nine opinions in which both findings were made, thirty-four (69%) of them found that all the factors favored a likelihood of confusion.”).

234. Michael Grynberg, *The Judicial Role in Trademark Law*, 52 B.C. L. REV. 1283, 1305 (2011) [hereinafter Grynberg, *Judicial Role in Trademark Law*].

235. Beebe, *supra* note 6, at 1614 (“Like any human decision makers, district judges attempt to decide both efficiently and accurately. In pursuit of efficiency, they consider only a few factors. In pursuit of accuracy, they consider the most decisive factors.”).

236. *Eclipse Assocs. Ltd. v. Data Gen. Corp.*, 894 F.2d 1114, 1118 (9th Cir. 1990) (“These tests were not meant to be requirements or hoops that a district court need jump through to make the determination.”); see also *R.H. Donnelley Inc. v. USA Northland Directories, Inc.*, No. Civ.04-4144, 2004 WL 2713248, at *6 (D. Minn. Nov. 19, 2004) (folding similarity and intent); *Ironhawk Techs. v. Dropbox, Inc.*, 994 F.3d 1107, 1123–24 (9th Cir. 2021) (folding actual confusion and sophistication); CDOC,

“core factors” informing “consumer perception in the marketplace rather than judicial perception in the courtroom.”²³⁷ Alejandro Mejías went further, recommending just two—similarity of marks and proximity of goods, as “adding any other relevant factors, instead of using unmanageable and misleading large lists of factors that are extremely difficult to balance, seems to be more in line with the thesis of scientific research on decision-making.”²³⁸ The next Part explains why actual confusion, mark similarity, and competitive proximity should form the core factors and why these factors, together with fair use safe harbors for expressive and descriptive uses, should form the rules of thumb in trademark law.

II. RULES OF THUMB

Mark similarity, goods and services, and evidence of actual confusion anchor the likelihood of confusion test as the most relevant factors.²³⁹ Jurisprudence supports that view. In one case from the dataset, the Ninth Circuit has described a “trinity [that] constitutes the most crucial body of the *Sleekcraft* analysis”—mark similarity, goods/services similarity, and marketing and advertising channels.²⁴⁰ Safe harbors protect core policies most in danger of being invaded by trademark expansionism while making it simpler and cheaper for businesses to do their due diligence and comply with the law.²⁴¹ This Part explains why.

A. Actual Confusion

Actual confusion is the most direct and decisive evidence of confusion.²⁴² Courts explain that where confusion occurred, it “is of

Inc. v. Liberty Bankers Life Ins. Co., 844 F. App’x 357, 362 (Fed. Cir. 2021) (folding competitive proximity and sophistication).

237. Beebe, *supra* note 6, at 1646.

238. Mejías, *supra* note 106, at 348 (concentrating the analysis on the main two factors).

239. See *infra* Part III.

240. GoTo.com, Inc. v. Walt Disney Co., 202 F.3d 1199, 1207 (9th Cir. 2000); see also Stone Creek, Inc. v. Omnia Italian Design, Inc., 862 F.3d 1131, 1136 (9th Cir. 2017) (“Two particularly probative factors are the similarity of the marks and the proximity of the goods.” (citing *Lindy Pen Co. v. Bic Pen Corp.*, 796 F.2d 254, 256–57 (9th Cir. 1986))).

241. See *infra* Part II.

242. Beebe, *supra* note 6, at 1608 (finding a 92% plaintiff success rate where the court found actual confusion); *Tana v. Dantanna’s*, 611 F.3d 767, 779 (11th Cir. 2010) (“The last factor, actual confusion in the consuming public, is the most persuasive evidence in assessing likelihood of confusion.”); see also Groeneveld Transp. Efficiency,

course convincing evidence that confusion is likely to occur.”²⁴³ As a policy lever, actual confusion gives courts the ability to anchor their analysis in real-world characteristics. In addition, the evidence is pre-existing, does not depend on the vagaries of survey design, and should make it easier for courts to dispose of cases pretrial.²⁴⁴


The dataset reveals that actual confusion comes up in 74% of the cases. In 32% of all cases, the courts favored plaintiffs on the actual confusion factor. When they did, plaintiffs won 54% of the time. In 18% of all cases, the courts favored defendants on the consumer sophistication factor. When the court favored defendants, they won 77% of the time.

Inc. v. Lubecore Int'l, Inc., 730 F.3d 494, 517 (6th Cir. 2013) (“Nothing shows the likelihood of confusion more than the fact of actual confusion.”); Variety Stores, Inc. v. Walmart Inc., 852 F. App'x 711, 720 (4th Cir. 2021) (“[A]ctual confusion, is the ‘most important factor’”); John Benton Russell, *New Tenth Circuit’s Standards: Competitive Keyword Advertising and Initial Interest Confusion in 1-800 Contacts v. Lens.com*, 30 BERKELEY TECH. L.J. 993, 1000 (2015) (“[C]ourts across several circuits view this as the strongest evidence a plaintiff can present in a trademark infringement case.”); Mark D. Robins, *Actual Confusion in Trademark Infringement Litigation: Restraining Subjectivity Through a Factor-Based Approach to Valuing Evidence*, 2 NW. J. TECH. & INTELL. PROP. 117, 117 (2004) (“In a case where all other circumstances point to a finding of non-infringement, significant evidence of actual confusion dramatically alters the equation.”).

243. *Morningside Grp. Ltd. v. Morningside Cap. Grp., L.L.C.*, 182 F.3d 133, 141 (2d Cir. 1999).

244. I am grateful to Jon Lee for this insight.

Figure 9: Actual Confusion by Outcome



At the bottom, the inquiry concerns whether there was confusion that could lead to “a diversion of sales, damage to goodwill, or loss of control over reputation.”²⁴⁵ For that reason, courts look for actual confusion among “prospective purchasers of [plaintiff’s] products.”²⁴⁶ Relevant circumstances include the extent of the parties’ advertising, the length of time the allegedly infringing product has been advertised, or any other factor that might influence the likelihood that actual confusion would be reported.²⁴⁷

Courts accept both anecdotal and survey evidence indicating actual confusion.²⁴⁸ This Article explained in Section I.C.2 that courts should avoid survey evidence in its current manifestation. As to anecdotal evidence, there is no absolute number of instances of actual confusion that must be met to win in a likelihood of confusion analysis. Rather, courts look to the totality of the circumstances in evaluating the evidence of actual confusion.²⁴⁹ For example, “[i]nquiries about the relationship between an owner of a mark and an alleged infringer do

245. Reply All Corp. v. Gimlet Media, LLC, 843 F. App’x 392, 397 (2d Cir. 2021) (quoting Lang v. Ret. Living Publ’g Co., 949 F.2d 576, 583 (2d Cir. 1991)).

246. Lang, 949 F.2d at 583; SLY Mag., LLC v. Weider Publ’ns L.L.C., 529 F. Supp. 2d 425, 441 (S.D.N.Y. 2007), *aff’d*, 346 F. App’x 721 (2d Cir. 2009).

247. Jellibeans, Inc. v. Skating Clubs Ga., Inc., 716 F.2d 833, 844 (11th Cir. 1983).

248. George & Co. v. Imagination Ent. Ltd., 575 F.3d 383, 398 (4th Cir. 2009).

249. AmBrit, Inc. v. Kraft, Inc., 812 F.2d 1531, 1543 (11th Cir. 1986).

not amount to actual confusion.²⁵⁰ Testimony from one customer—the mark’s owner, and its employee—arguing customers mistakenly visited the defendant’s store when intending to visit the owner’s store constitutes *de minimis* evidence of actual confusion.²⁵¹ While “[i]solated instances of [actual] confusion are insufficient to support a finding of likely confusion,”²⁵² courts have found confusion by five people,²⁵³ or even one person increases the likelihood of confusion.²⁵⁴ At the same time, “it is well established that no actual confusion is required to prove a case of trademark infringement.”²⁵⁵ Courts have justified this conclusion “[b]ecause of the difficulty in garnering such evidence.”²⁵⁶

Confusion must be by the “actual consuming public” and therefore anchored in a real-world context.²⁵⁷ The absence of actual confusion “over a substantial period . . . creates a strong inference that there is no likelihood of confusion.”²⁵⁸ In combining the two ideas, “[s]hort-lived confusion or confusion of individuals casually acquainted with a business is worthy of little weight, while confusion of actual customers of a business is worthy of substantial weight.”²⁵⁹ Where a large volume of contacts or transactions could give rise to confusion, and only limited instances of confusion present themselves, courts give evidence of actual confusion little weight.²⁶⁰

250. *Reply All Corp.*, 843 F. App’x at 398 (alteration in original).

251. *RVC Floor Decor, Ltd. v. Floor & Decor Outlets of Am., Inc.*, 527 F. Supp. 3d 305, 326 (E.D.N.Y. 2021).

252. *Progressive Distrib. Servs., Inc. v. United Parcel Serv., Inc.*, 856 F.3d 416, 433 (6th Cir. 2017).

253. *AWGI, LLC v. Atlas Trucking Co.*, 998 F.3d 258, 267 (6th Cir. 2021).

254. *Innovation Ventures, LLC v. N2G Distrib., Inc.*, 763 F.3d 524, 536 (6th Cir. 2014) (“[A] single instance of actual confusion can, in some cases, ‘increase the likelihood of confusion.’”).

255. *Louis Vuitton Malletier S.A. v. Haute Diggity Dog, LLC*, 507 F.3d 252, 263 (4th Cir. 2007).

256. *Monster Energy Co. v. BeastUp LLC*, 395 F. Supp. 3d 1334, 1358 (E.D. Cal. 2019) (alteration in original).

257. *Rearden LLC v. Rearden Com., Inc.*, 683 F.3d 1190, 1210 (9th Cir. 2012).

258. *CareFirst of Md., Inc. v. First Care, P.C.*, 434 F.3d 263, 269 (4th Cir. 2006) (finding an inference of no likelihood of confusion where there was no evidence of confusion for nine years).

259. *Safeway Stores, Inc. v. Safeway Disc. Drugs, Inc.*, 675 F.2d 1160, 1167 (11th Cir. 1982).

260. *George & Co. v. Imagination Ent. Ltd.*, 575 F.3d 383, 399 (4th Cir. 2009) (explaining that “the company’s failure to uncover more than a few instances of actual confusion creates a presumption against likelihood of confusion in the future” when there are so many opportunities for confusion to occur).

Some courts hold that a lack of evidence of actual confusion does not create a presumption of no confusion but is “simply a factor in the court’s analysis.”²⁶¹ However, as a policy lever, it gives courts the ability to anchor their analysis in real-world characteristics. In addition, the evidence is pre-existing, does not depend on the vagaries of survey design, and should make it easier for courts to dispose of cases pretrial.²⁶² If found, it is worth its weight in gold, tipping the balance in the plaintiff’s favor more than any other factor.²⁶³

B. *Mark Similarity*

Three axioms apply to the “similarity” analysis: (1) marks should be considered in their entirety and as they appear in the marketplace; (2) similarity is judged by appearance, sound, and meaning; (3) and, similarities weigh more heavily than differences.²⁶⁴ Courts determine whether a mark confuses the public when viewed alone to account for the possibility that similar marks “may confuse consumers who do not have both marks before them but who may have a general, vague, or even hazy, impression or recollection of the other party’s mark.”²⁶⁵

At the most basic level, marks are confusingly similar if “ordinary consumers would likely conclude that [the two products] share a common source, affiliation, connection or sponsorship.”²⁶⁶ Identical, even dominant, features do not “automatically mean that two marks are similar.”²⁶⁷ Courts look to “the overall impression created by the marks, not merely compare individual features,” and “may consider the marks’ visual, aural, and definitional attributes and compare the trade dress of the products in determining whether the total effect conveyed by the two marks is confusingly similar.”²⁶⁸

261. See, e.g., *Am. Dairy Queen Corp. v. W.B. Mason Co.*, 543 F. Supp. 3d 695, 718–19 (D. Minn. 2021).

262. I am grateful to Jon Lee for this insight.

263. “Bridging the gap” reported a 71% win rate but given its relative infrequency (25% versus 74% for actual) and large overlap with competitive proximity (which could explain why it is even at 25%, for that matter 71%), the better view is to discount it. *Lim*, *supra* note 41.

264. *Entrepreneur Media, Inc. v. Smith*, 279 F.3d 1135, 1144 (9th Cir. 2002).

265. *Maker’s Mark Distillery, Inc. v. Diageo N. Am., Inc.*, 679 F.3d 410, 421 (6th Cir. 2012).

266. *Fisons Horticulture, Inc. v. Vigoro Indus., Inc.*, 30 F.3d 466, 477 (3d Cir. 1994).

267. *Sensient Techs. Corp. v. SensoryEffects Flavor Co.*, 613 F.3d 754, 764 (8th Cir. 2010) (quoting *Gen. Mills, Inc. v. Kellogg Co.*, 824 F.2d 622, 627 (8th Cir. 1987)).

268. *Luigino’s, Inc. v. Stouffer Corp.*, 170 F.3d 827, 830 (8th Cir. 1999); see also *Leelanau Wine Cellars, Ltd. v. Black & Red, Inc.*, 502 F.3d 504, 516–17 (6th Cir. 2007) (explaining that courts give similarity considerable weight and consider the

As the Eleventh Circuit noted, “[i]f a trademark operates in a crowded field of similar marks on similar goods or services, slight differences in names may be meaningful because consumers will not likely be confused between any two of the crowd and may have learned to carefully pick out one from the other.”²⁶⁹ Similarly, mark similarity takes prominence when the goods are direct competitors in the marketplace.²⁷⁰ The goods or services will likely be virtually identical. However, the marks need not be as similar for there to be a likelihood of confusion.²⁷¹

Courts extol the importance of mark similarity.²⁷² Beebe found it was “by far the most important factor.”²⁷³ In injunction cases, 83% of plaintiffs who won the similarity factor prevailed in the likelihood of confusion analysis, as did 90% in plaintiff summary judgment motions.²⁷⁴ The dataset reveals that only 23% of plaintiffs in injunction cases who won the similarity factor prevailed in likelihood of confusion analysis, a far lower number than before. However, 90% in plaintiff summary judgment motions succeeded just as before.

Fifteen years later, the dataset also reveals that mark similarity still comes up in 85% of the cases, the most frequently invoked factor of them all. In 61% of all cases, the courts favored plaintiffs on the mark similarity factor. When they did, plaintiffs won 47% of the time on the merits. In 18% of all cases, the courts favored defendants on the mark similarity factor. When the court favored defendants, they won 88% of the time on the merits.

“pronunciation, appearance, and verbal translation of conflicting marks,” and courts must “view marks in their entirety and focus on their overall impressions, not individual features” (quoting *AutoZone, Inc. v. Tandy Corp.*, 373 F.3d 786, 795–96 (6th Cir. 2004)).

269. *Fla. Int’l Univ. Bd. Trs. v. Fla. Nat’l Univ., Inc.*, 830 F.3d 1242, 1260 (11th Cir. 2016) (internal quotation marks omitted).

270. *Bos. Duck Tours, LP v. Super Duck Tours, LLC*, 531 F.3d 1, 30 (1st Cir. 2008).


271. *Bridgestone Ams. Tire Operations, LLC v. Fed. Corp.*, 673 F.3d 1330, 1337 (Fed. Cir. 2012) (likelihood of confusion between “Potenza” and “Turanza” marks was greater because both referred to tires).

272. *Equitable Nat’l Life Ins. Co. v. AXA Equitable Life Ins. Co.*, 434 F. Supp. 3d 1227, 1246 (D. Utah 2020) (“The similarity of the marks is the ‘first and most important factor.’”); *Juul Labs, Inc. v. 4X PODS.*, 439 F. Supp. 3d 341, 353 (D.N.J. 2020) (calling it “[t]he single most important factor in determining likelihood of confusion”); *GoTo.com, Inc. v. Walt Disney Co.*, 202 F.3d 1199, 1205 (9th Cir. 2000) (noting that mark similarity “has always been considered a critical question in the likelihood-of-confusion analysis”).

273. Beebe, *supra* note 6, at 1623.

274. *Id.* at 1625.

Figure 10: Mark Similarity by Outcome



The dataset reveals interesting dynamics between rivalry and mark similarity. When the parties' marks were similar and the parties were rivals, plaintiffs won 47% of the time. Similarly, when the parties were non-rivals, plaintiffs also won 47% of the time. However, there is a difference in how often the defendant wins. Defendants only win 10% of the time when they are rivals and 24% when they are not rivals, underscoring the expectedly powerful role rivalry plays, but in an asymmetrical way.

One possible explanation is that similarity between the marks makes consumers more likely to become confused about the source. Extremely similar marks or goods may suggest counterfeiting and free riding. Parodies, comparative advertising, and nominative use make consumers less likely to be confused, even if the third party uses the identical term.

Courts even dispensed entirely with the likelihood of confusion test when parties' marks were identical, a conclusion with implications for early off-ramping parties, as discussed in Part III.²⁷⁵ Where a defendant uses a counterfeit mark, such use is deemed inherently confusing to a

275. See *Phillip Morris USA Inc. v. Shalabi*, 352 F. Supp. 2d 1067, 1073 (C.D. Cal. 2004) (“[I]n cases involving counterfeit marks, it is unnecessary to perform the step-by-step examination . . . because counterfeit marks are inherently confusing.”); *Daimler AG v. A-Z Wheels LLC*, 334 F. Supp. 3d 1087, 1096 (S.D. Cal. 2018) (“It is not necessary for the Court to analyze the likelihood of confusion test here considering Defendants’ use the identical MERCEDES-BENZ mark.”).

customer.²⁷⁶ As McCarthy explained “[c]ases where a defendant uses an identical mark on competitive goods hardly ever find their way into the appellate reports. Such cases are ‘open and shut’ and do not involve protracted litigation to determine liability for trademark infringement.”²⁷⁷ This is because “confusing the customer is the whole purpose of creating counterfeit goods.”²⁷⁸ Such cases create a presumption of harm such that the factor may stampede the likelihood of confusion analysis entirely.²⁷⁹

Aside from the simplest forms of counterfeiting, the threshold triggering confusion, and more so *likely* confusion, exists only as a relative measure where reasonable minds may differ. Unlike real property, there are no metes and bounds. This lack of boundaries presents interpretive challenges that Michael Grynberg and Graeme Austin independently attributed to the likelihood of confusion’s current uncertainty.²⁸⁰ The problem is common to other areas of the law as well. For instance, copyright law’s substantial similarity standard suffers many of the same ills as the likelihood of confusion and demands reconsideration.²⁸¹

276. *EAT BBQ LLC v. Walters*, 47 F. Supp. 3d 521, 530 (E.D. Ky. 2014) (“[T]here is almost never a dispute regarding confusion.”).

277. *Wynn Oil Co. v. Thomas*, 839 F.2d 1183, 1191 (6th Cir. 1988) (quoting 4 MCCARTHY, *supra* note 2, § 23:20).

278. *Gucci Am., Inc. v. Duty Free Apparel, Ltd.*, 286 F. Supp. 2d 284, 287 (S.D.N.Y. 2003).

279. *See Threeline Imps., Inc. v. Vernikov*, 239 F. Supp. 3d 542, 561 (E.D.N.Y. 2017); *see also Dish Network L.L.C. v. Siddiqi*, No. 18 CV 4397, 2019 WL 5781945, at *4 (S.D.N.Y. 2019) (“[I]t is not necessary to perform the step-by-step examination of each *Polaroid* factor’ when a counterfeit mark is at issue.”); *Halo Optical Prods., Inc. v. Liberty Sport, Inc.*, No. 14-cv-00282, 2017 WL 1082443, at *11 (N.D.N.Y. Mar. 22, 2017) (“[W]hen dealing with an identical mark . . . courts are not necessarily required to analyze the *Polaroid* factors.”); *Gucci Am.*, 286 F. Supp. 2d at 287 (“[T]he Court need not undertake a factor-by-factor analysis under *Polaroid* because counterfeits, by their very nature, cause confusion.”).

280. Grynberg, *Judicial Role in Trademark Law*, *supra* note 234, at 1303 (“Trademark’s fundamental inquiry, whether a likelihood of confusion exists, invites judicial lawmaking in no small part because the term ‘likelihood of confusion’ presents an interpretive problem.”); Graeme W. Austin, *Tolerating Confusion About Confusion: Trademark Policies and Fair Use*, 50 ARIZ. L. REV. 157, 160 (2008) (“There is considerable uncertainty about some of the key questions that are germane to the factual inquiry at the heart of the likelihood of confusion analysis.”).

281. *See, e.g., Daryl Lim, Saving Substantial Similarity*, 73 FLA. L. REV. 591, 640–41 (2021) [hereinafter *Lim, Saving Substantial Similarity*] (explaining how the substantial similarity standard generates “capricious and wrong results”).

Courts use sights, sounds, and meaning to make snap judgments about mark similarity.²⁸² These heuristics allow judges to rely on “a small set of cheap and reliable factors that are close enough to the ideal.”²⁸³ Adam Samaha approves of it since “[p]rioritizing the judge’s impressions about the similarity of marks, therefore, tends toward the high values of trademark law at bargain basement prices.”²⁸⁴ Defendants can easily compare visual or aural elements in context, making this a useful factor to encourage due diligence.²⁸⁵

The key takeaway is that the commercial context matters.²⁸⁶ Marks should not be compared side-by-side as they might be shown in the courtroom.²⁸⁷ Instead, courts determine whether the public would confuse the marks when viewed alone because some highly similar marks can confuse consumers that view them without appropriate commercial context.²⁸⁸ For this reason, courts cannot dissect marks since consumers encounter them in their entirety in those settings. Instead, courts focus on their overall impressions rather than on their features.²⁸⁹ That which qualifies as mark similarity also disqualifies intent, survey evidence, mark strength, and consumer sophistication.

C. *Competitive Proximity*

Competitive proximity tells courts how likely consumers are to assume an association between the marks used on related products.²⁹⁰ For example, “[t]he similarities between the parties’ distribution channels and marketing strategies suggest an overlapping general class

282. Adam M. Samaha, *Looking Over a Crowd—Do More Interpretive Sources Mean More Discretion?*, 92 N.Y.U. L. REV. 554, 614 (2017) (“[A]ccurately estimating the probability of consumer confusion can require a snap judgment, which often is how consumers actually formulate impressions and make purchasing decisions.”).

283. *Id.*

284. *Id.*

285. See 4 MCCARTHY, *supra* note 2, § 23:21 (discussing the “sound, sight, and meaning” test for mark similarity).

286. *Flower Mfg., LLC v. CareCo, LLC*, 466 F. Supp. 3d 797, 814 (N.D. Ohio 2020) (“I am to judge the marks’ similarity as they appear in their commercial context.”); see also *Homeowners Grp., Inc. v. Home Mktg. Specialists, Inc.*, 931 F.2d 1100, 1109 (6th Cir. 1991) (“[A] court must determine, in the light of what occurs in the marketplace, whether the mark will be confusing to the public when singly presented.” (internal quotation marks omitted) (quoting *Wynn Oil Co. v. Thomas*, 839 F.2d 1183, 1187 (6th Cir. 1988))).

287. *Homeowners Grp., Inc.*, 931 F.2d at 1106.

288. *Daddy’s Junky Music Stores, Inc. v. Big Daddy’s Fam. Music Ctr.*, 109 F.3d 275, 283 (6th Cir. 1997) (quoting *Wynn Oil*, 839 F.2d at 1188).

289. *AutoZone, Inc. v. Tandy Corp.*, 373 F.3d 786, 795 (6th Cir. 2004).

290. *Virgin Enters. Ltd. v. Nawab*, 335 F.3d 141, 150 (2d Cir. 2003).

of consumers of the parties' products."²⁹¹ However, two products or services within the same general field do not automatically trigger a likelihood of confusion.²⁹² Similarly, a high percentage of overlap in "an extremely small subset of products does not demonstrate a high degree of relatedness."²⁹³

Services and goods within the same broad industry are not necessarily related. Rather, related services are marketed and consumed such that buyers are likely to believe that the services come from a common company.²⁹⁴ Courts examine "how and to whom the respective goods or services of the parties are sold."²⁹⁵ Less likelihood of confusion exists where the goods are sold through different avenues, "parties have different customers[,] and [they] market their goods or services in different ways."²⁹⁶ "[I]f the parties compete directly, confusion is likely" between sufficiently similar marks.²⁹⁷ "[I]f the goods and services are somewhat related, but not competitive, then the likelihood of confusion will turn on other factors[.]"²⁹⁸ [F]inally, if the products are unrelated, confusion is highly unlikely.²⁹⁹

The dataset reveals that competitive proximity comes up in 74% of the cases. In 51% of all cases, the courts favored plaintiffs on the competitive proximity factor. When they did, plaintiffs won 45% of the time. In 14% of all cases, the courts favored defendants on the competitive proximity factor. When the court favored defendants, they won 85% of the time.

291. *Monster Energy Co. v. BeastUp LLC*, 395 F. Supp. 3d 1334, 1359 (E.D. Cal. 2019). Other circuits use similar formulations. *See e.g.*, *Therma-Scan, Inc. v. Thermoscan, Inc.*, 295 F.3d 623, 632 (6th Cir. 2002) (noting that direct rivalry through similar goods or services is likely confusing).

292. *Matrix Motor Co. v. Toyota Jidosha Kabushiki Kaisha*, 290 F. Supp. 2d 1083, 1092 (C.D. Cal. 2003).

293. *AutoZone, Inc.*, 373 F.3d at 798 ("[I]f [the defendant] stocked only five types of batteries all of which were also sold by [the plaintiff], the overlap would be 100%, even though in reality [the defendant] and [the plaintiff] would share only five products of the approximately 55,000 offered by [the plaintiff].").

294. *Daddy's Junky Music Stores, Inc. v. Big Daddy's Fam. Music Ctr.*, 109 F.3d 275, 282–83 (6th Cir. 1997) (quoting *Homeowners Grp., Inc. v. Home Mktg. Specialists, Inc.*, 931 F.2d 1100, 1109 (6th Cir. 1991)).

295. *Leelanau Wine Cellars, Ltd. v. Black & Red, Inc.*, 502 F.3d 504, 519 (6th Cir. 2007).

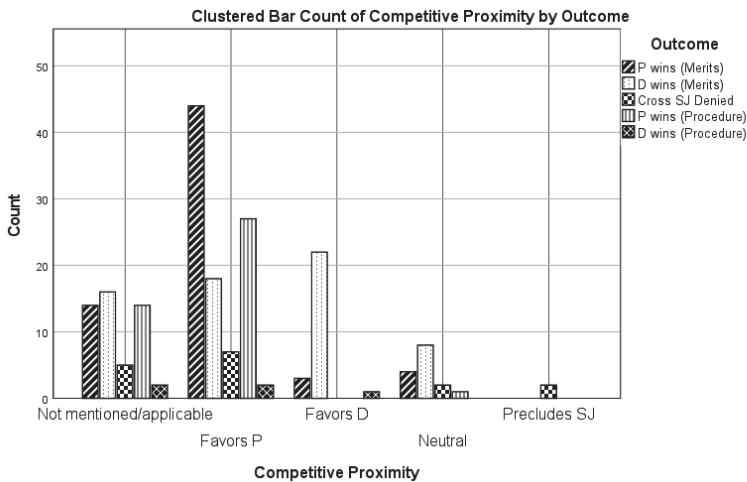
296. *Therma-Scan*, 295 F.3d at 636.

297. *Kellogg Co. v. Toucan Golf, Inc.*, 337 F.3d 616, 624 (6th Cir. 2003).

298. *Id.*

299. *Id.*

Figure 11: Competitive Proximity by Outcome



Competitive proximity encompasses adjacent *Polaroid* factors. One example is the likelihood plaintiffs or defendants will expand into each other's market or the "bridging the gap" factor.³⁰⁰ The likelihood that consumers will confuse the sources of parties' products increases when there is a "strong possibility that either party will expand its business to compete with the other's."³⁰¹ This confusion may happen when goods and services are complementary, sold to the same class of purchasers, or similar in use and function.³⁰²

Courts examine the two concepts in tandem with each other.³⁰³ Consider *Kohler Co. v. Bold International FZCO*,³⁰⁴ where the court noted that "[b]ridging the gap' refers to the likelihood that the senior

300. *Disney Enters., Inc. v. Sarelli*, 322 F. Supp. 3d 413, 434 (S.D.N.Y. 2018) ("The third and fourth *Polaroid* factors, respectively, address the proximity of the goods or services at issue and the possibility that the senior user will 'bridge the gap,' or expand the scope of its business and enter the market of the junior user. Thus, these two distinct but related factors 'focus on the degree to which the [parties'] products currently compete with each other or are likely to compete with each other in the future.'" (alteration in original) (quoting *Medici Classics Prods., LLC v. Medici Grp., LLC*, 683 F. Supp. 2d 304, 311–12 (S.D.N.Y. 2010))).

301. *Kibler v. Hall*, 843 F.3d 1068, 1082 (6th Cir. 2016).

302. *AMF Inc. v. Sleekcraft Boats*, 599 F.2d 341, 350 (9th Cir. 1979).

303. *RVC Floor Decor, Ltd. v. Floor & Decor Outlets Am., Inc.*, 527 F. Supp. 3d 305, 325 (E.D.N.Y. 2021) ("For the same reasons explained [in] the 'competitive proximity' analysis, the parties serve the same market and any gap has already been bridged.").

304. 422 F. Supp. 3d 681 (E.D.N.Y. 2018).

user . . . will enter into the same market as that of the junior user . . . where the goods are not yet in close competitive proximity.”³⁰⁵ When the parties’ goods are the same, courts simply fold this factor into competitive proximity as there is no gap to bridge.³⁰⁶ In this case, a consumer seeing the goods or services would likely be confused about their source.³⁰⁷

Another example is the degree of care the consumer might exercise in purchasing the parties’ goods, as mentioned in Section I.C.4. Courts look both to the “relative sophistication of the relevant consumer”³⁰⁸ and the cost of the item³⁰⁹ in determining the degree of care likely to be exercised by the purchaser. The “reasonably prudent consumer” is expected “to be more discerning—and less easily confused—when [they are] purchasing expensive items.”³¹⁰ Conversely, customers may be less careful when purchasing inexpensive products, thus making confusion more likely.³¹¹

“Bridging the gap” rarely arose, only in 25% of the cases. In 7% of all cases, the courts favored plaintiffs on the “bridging the gap” factor. When they did, plaintiffs won 71% of the time. In 6% of all cases, the courts favored defendants on the “bridging the gap” factor. When the court favored defendants, they won 83% of the time.

305. *Id.* at 725.

306. *Star Indus., Inc. v. Bacardi & Co.*, 412 F.3d 373, 387 (2d Cir. 2005) (“Because . . . [the parties’] products are already in competitive proximity, there is really no gap to bridge, and this factor is irrelevant to the *Polaroid* analysis in this case.”).

307. *DJ Direct, Inc. v. Margalot*, 512 F. Supp. 3d 396, 411 (E.D.N.Y. 2021).


308. *Fortune Dynamic, Inc. v. Victoria’s Secret Stores Brand Mgmt., Inc.*, 618 F.3d 1025, 1038 (9th Cir. 2010).

309. *Brookfield Commc’ns, Inc. v. W. Coast Ent. Corp.*, 174 F.3d 1036, 1060 (9th Cir. 1999).

310. *Id.*


311. *Id.*

Figure 12: "Bridging the Gap" by Outcome



Similarly, the quality of the defendant’s goods is closely related to competitive proximity. The quality factor featured even more infrequently, and 11% less than “bridging the gap.” In 1% of all cases, the courts favored plaintiffs on the quality factor. When they did, plaintiffs won 50% of the time. In 1% of all cases, the courts favored defendants on the “bridging the gap” factor. When the court favored defendants, they won 100% of the time. Of course, these figures should be seen in the context of the very small sample size.

Figure 13: Quality of D's Product by Outcome



Courts look to the product, the relevant market, and potential consumers.³¹² Product proximity overlaps substantially with marketing and advertising channels and should be subsumed within those channels. For this reason, product proximity can serve as an omnibus factor for other factors such as the relative quality of goods sold, “bridging the gap” from the perspective of the relevant public (rather than from the legitimate aspirations of the trademark owner), and similarity of distribution channels.

One court in the dataset was exemplary in defining the relevant consumer market.³¹³ To determine whether that market included potential commercial and government customers, it examined the trademark owner’s revenue sources, proposals it sent to two potential customers, and the defendant’s exploratory acquisition of the trademark owner to conclude that both parties targeted similar customers.³¹⁴ Despite this fact-intensive inquiry, the court notably

312. *Best Cellars, Inc. v. Grape Finds at Dupont, Inc.*, 90 F. Supp. 2d 431, 456 (S.D.N.Y. 2000).

313. *Ironhawk Techs., Inc. v. Dropbox, Inc.*, 994 F.3d 1107, 1117 (9th Cir. 2021) (“Before addressing the *Sleekcraft* factors, we must define the relevant consumer market because ‘a court conducting a trademark analysis should focus its attention on the relevant consuming public.’” (quoting *Rearden LLC v. Rearden Com., Inc.*, 683 F.3d 1190, 1214 (9th Cir. 2012))).

314. *Id.* at 1117–18.

concluded that “a reasonable jury could find that [plaintiff’s] potential consumers include commercial customers.”³¹⁵

Market definition was an interesting issue that arose in a few cases in the dataset. Though similar, courts distinguish between market definition under trademark and antitrust law. For example, one case in the dataset reported a plaintiff asserting that “courts have looked to antitrust law . . . to find goods competitive where they are ‘either identical or available substitutes for each other.’”³¹⁶ Disagreeing with this assessment, the court responded that “the question . . . is not whether [the defendant’s] conduct impair[ed] competition in the marketplace[,] but whether it . . . infringed” upon a protected interest in the plaintiff’s trademark.³¹⁷

D. *Summing It Up*

In sum, the eight *Polaroid* Factors can be efficiently subsumed into a troika of actual confusion, mark similarity, and competitive proximity. The table below shows the troika being the most prominent factors. They also deliver consistent win rates to plaintiffs if the particular mark favors them, at between 45% to 54%, mapping almost exactly to Priest-Klein’s 50% figure discussed in Section I.A.

315. *Id.* at 1118.

316. *Flower Mfg., LLC v. CareCo, LLC*, 466 F. Supp. 3d 797, 812 (N.D. Ohio 2020) (emphasis omitted).

317. *Id.* (“It is fair to say that trademark laws were enacted for the protection of the competitor who owns a mark and not for protection of competition in the marketplace in general.”).

Figure 14: Revised Likelihood of Confusion Factors

<i>Polaroid</i> Factors	Rules of Thumb	Frequency (%)	Plaintiff's Win Rate (%)
Strength of the plaintiff's mark	Discarded	70	46
Similarity of plaintiff's and defendant's marks	Retained	85	47
Competitive proximity of products or services	Retained	73	45
Likelihood that plaintiff will "bridge the gap" and offer a product like a defendant's	Covered by competitive proximity of products or services	25	71
Actual confusion	Retained	74	54
Defendant's good faith	Discarded	66	52
Quality of defendant's product	Covered by competitive proximity of products or services	11	50
Buyer sophistication	Discarded	46	49

More importantly, the troika moves trademark doctrine a step in the right direction by limiting ad hoc fact-finding. However, the troika alone is incomplete. Mark McKenna and Mark Lemley warn that unless we can "identify more specifically the types of relationships that could give rise to actionable confusion, there is no logical stopping point for trademark protection."³¹⁸ The converse is also true—we also need to identify safe harbors. It is difficult even for savvy parties to predict the outcome in advance and resolve disputes early in proceedings, placing swathes of activity at significant risk.³¹⁹

318. See Mark A. Lemley & Mark McKenna, *Irrelevant Confusion*, 62 STAN. L. REV. 413, 439 (2010).

319. David S. Welkowitz, *The Virtues and Vices of Clarity in Trademark Law*, 81 TENN. L. REV. 145, 148 (2013) ("Because the level and even the existence of confusion is difficult to predict in advance, partly due to the uncertainties built into trademark law's test for confusion, those who would engage in valued activity must do so at significant risk.").

Simplifying confusion benefits other aspects of trademark law. For example, trademark law's first sale doctrine also permits some marketplace confusion by letting others sell used or reconditioned goods bearing the mark.³²⁰ Nominative fair use may likewise fold the likelihood of confusion standard into its analysis.³²¹ What is "fair" implicates the confusion arising from using the offending mark—whether the defendant only used as much as necessary of the plaintiff's mark—which in turn impacts the vagueness of the likelihood of confusion standard.³²² The same issue arises with expressive trademark uses³²³ and the legality of keyword advertising.³²⁴ Fair use is the focus of the next Section.

E. A Word on Fair Use

As trademarks expand beyond source identification, they seed public discourse with their communicative value.³²⁵ Trademark owners obtain rights with inchoate boundaries. When the public interacts with a trademark, the mark may imbue with collective meaning. This collective meaning has social value, and in appropriate instances, the law should offer them categorical protection from lawsuits.³²⁶

320. See *Nitro Leisure Prods., L.L.C. v. Acushnet Co.*, 341 F.3d 1356, 1362–64 (Fed. Cir. 2003) (consumer confusion as benchmark for applying the first sale doctrine).

321. *Toyota Motor Sales, U.S.A., Inc. v. Tabari*, 610 F.3d 1171, 1175–76 (9th Cir. 2010) (asking whether (1) the product was readily identifiable without use of the mark; (2) defendant used more of the mark than necessary; or (3) defendant falsely suggested he was sponsored or endorsed by the trademark holder).

322. *E.g.*, *KP Permanent Make-Up, Inc. v. Lasting Impression I, Inc.*, 543 U.S. 111, 123 (2004) (confusion relevant to whether descriptive use is "fair").

323. See *Rogers v. Grimaldi*, 875 F.2d 994, 999 (2d Cir. 1989) (adopting balancing test that asks whether the use of a trademark as the title of an expressive work is artistically relevant to the underlying work and, if so, whether "the title explicitly misleads as to the source or the content of the work").

324. See *Network Automation, Inc. v. Advanced Sys. Concepts, Inc.*, 638 F.3d 1137, 1154 (9th Cir. 2011) (noting that keyword advertisements could be "confusingly labeled or not labeled at all" making how advertisements appear on the results page must be considered).

325. See Alex Kozinski, *Trademarks Unplugged*, 68 N.Y.U. L. REV. 960, 973–74 (1993) (noting how businesses inject the "effervescent qualities" of trademarks "into the stream of communication with the pressure of a firehose by means of mass media campaigns").

326. See, *e.g.*, William McGeeveran & Mark P. McKenna, *Confusion Isn't Everything*, 89 NOTRE DAME L. REV. 253, 301–06 (2013) (proposing categorical exclusions for some favored uses).

Communication relies on a plethora of legally protected words, graphics, sounds, and smells.³²⁷ Beyond computers or smartphones, APPLE may represent a nonconformist hip lifestyle compared with users of LENOVO's more staid business offerings. Trademarks become tools of communication and expression, and the public helps shape their boundaries as they become symbols that embody culture itself.³²⁸

Trademark owners may be anxious to protect themselves from uses that dilute the value of a household logo or name even when consumers are not confused. Between 2019 and 2021, Apple filed 215 trademark oppositions, targeting small companies and nonprofits that have nothing to do with providing technology products or services, including an Indian food blog and a public school.³²⁹ In these mass-produced, boilerplate-worded oppositions, Apple has argued that "Apple marks are so famous and instantly recognizable" that other trademarks will weaken the strength of its brand or cause the "ordinary consumer to believe that applicant is related to, affiliated with or endorsed by Apple."³³⁰ While Apple protested that this is simply what the law "requires," Professor Christine Farley has called them "bullying tactics."³³¹ Whatever the case, the impact is clear; Apple has been successful in preventing registration of a wide variety of marks. When faced with an opposition by Apple, applicants expressly withdrew their applications 17% of the time and failed to respond and subsequently defaulted 59% of the time.³³² The Tech Transparency Report noted only one win against Apple—by the U.S. government:

327. Diamond & Franklyn, *supra* note 165, at 2031.

328. Barton Beebe, *The Semiotic Analysis of Trademark Law*, 51 UCLAL. REV. 621, 624 (2004) (arguing trademark law is both an economic doctrine and "a semiotic doctrine elaborating the principles of sign systems, of language").

329. See *Apple's Trademark 'Bullying' Targets Small Businesses, Nonprofits* TECH TRANSPARENCY PROJECT (Mar. 11, 2022) [hereinafter TECH TRANSPARENCY PROJECT], <https://www.techtransparencyproject.org/articles/apples-trademark-bullying-targets-small-businesses-nonprofits> [<https://perma.cc/96QR-UZBT>] ("Many choose to simply give up rather than take on a mega-corporation with a market value of \$2.5 trillion, but in doing so, they lose whatever funds they invested in designing their logo and hiring a lawyer to deal with the trademark application.").

330. Ryan Mac & Kellen Browning, *Apps and Oranges: Behind Apple's 'Bullying' on Trademarks*, N.Y. TIMES (Mar. 11, 2022) <https://www.nytimes.com/2022/03/11/technology/apple-trademarks.html> [<https://perma.cc/LEV8-RPEE>].

331. *Id.* For an interesting discussion on how to prevent "bullying" in trademark law, see Stacey Dogan, *Bullying and Opportunism in Trademark and Right-of-Publicity Law*, 96 B.U. L. REV. 1293 (2016) (examining various instances of trademark bullying, including by Monster Energy, one of the most infamous trademark bullies).

332. See TECH TRANSPARENCY PROJECT, *supra* note 329 ("Of the 118 North American cases analyzed by TTP, 76 have been decided in Apple's favor, with a complete defeat

Apple appears to have been stopped in only one instance: when it challenged the U.S. government. The Department of Energy registered an online research service called Pages in 2020. Apple owns the trademark Pages for its word processor and opposed the agency's trademark on those grounds. The Department's response was brief, and its main defense took the form of a one-sentence declaration: "There is no likelihood of confusion, mistake, or deception between [Apple's] marks and DOE's PAGES mark." After more than a year of negotiations, Apple agreed to withdraw its objections, without DOE making any change to its application. In TTP's dataset, no other applicant was able to completely beat back Apple.³³³

While the costs of invading free speech and other interests are high, the costs of being overly permissive in expressive use cases cause only minimal harm.³³⁴ Research on brand extensions shows owners are rarely "harmed by consumers' mistaken association of unrelated products."³³⁵ Consumers rarely alter how they view the brand quality when they encounter criticism about other products offered under that mark.³³⁶ The negative impact stays with the related products but does not corrupt a positive view of the owner's line of products.³³⁷

For this reason, Mark McKenna has warned against reflexively prohibiting every form of confusion.³³⁸ Instead, trademark law should only be concerned with confusion that influences consumer decision-making. Bone also cautions that economic concerns over confusion should be distinguished from penalizing intentional deception without evidence of consumer confusion.³³⁹ Similarly, Lisa Ramsey flags the need to safeguard free speech interests in the face of

of the other parties' proposed logos or trademarks. None of these cases saw a full trial before TTAB but were withdrawn or abandoned by the applicants amid Apple's pressure.").

333. *Id.*

334. William McGeeveran, *The Trademark Fair Use Reform Act*, 90 B.U. L. REV. 2267, 2286 (2010).

335. Lemley & McKenna, *supra* note 318, at 429.

336. *See id.* at 429.

337. *See id.* at 430 ("Consumers, in other words, are smart enough to distinguish different products and hold different impressions of them.").

338. Mark P. McKenna, *A Consumer Decision-Making Theory of Trademark Law*, 98 VA. L. REV. 67, 73 (2012).

339. Robert G. Bone, *Taking the Confusion out of "Likelihood of Confusion": Toward a More Sensible Approach to Trademark Infringement*, 106 NW. L. REV. 1307, 1377 (2012).

encroaching trademark enforcement.³⁴⁰ One way to do this is by fortifying fair use.

Fair use is currently regarded as an affirmative defense.³⁴¹ That generally precludes the pretrial disposition of the case.³⁴² In the interest of early off-ramping cases, it is perhaps fortunate then that fair use rarely arose in the dataset (6%), with equal probability that a court would eventually find in favor of either plaintiff (42%) or defendant (50%). It also provides empirical evidence that converting fair use from an affirmative defense to a safe harbor would create a powerful tool to fend off trademark trolls without appreciably disrupting day-to-day trademark practice.

Safe harbors offer advantages over attempts to prescribe clear rules. These include improving predictability and ease of determination, allowing courts to resolve issues sooner in the litigation process. Here, Gideon Parchomovsky and Alex Stein make a more general point that “[r]eplacing these criteria with rules that will lay down irrebuttable presumptions of consumer confusion, or lack thereof, could make litigation over trademarks cheaper than it presently is.”³⁴³ The case is over as soon as the defendants demonstrate a basic fact.³⁴⁴

Safe harbors exist within trademark law, specifically the likelihood of confusion tests. For instance, the law does not protect functional product designs to avoid giving plaintiffs an advantage against rivals unrelated to the plaintiff’s reputation.³⁴⁵ Similarly, the law keeps plaintiffs on a leash to not monopolize trademarks with descriptive words and receive protection for generic terms.³⁴⁶ Expressive uses for

340. See Lisa P. Ramsey, *Increasing First Amendment Scrutiny of Trademark Law*, 61 SMU L. REV. 381, 384–85 (2008) (advocating rigorous First Amendment rights of trademark laws to protect free speech).

341. See, e.g., *KP Permanent Make-Up, Inc. v. Lasting Impression I, Inc.*, 543 U.S. 111, 123 (2004) (finding that confusion is relevant to whether descriptive use is “fair”); *New Kids on the Block v. News Am. Publ’g, Inc.*, 971 F.2d 302, 308 (9th Cir. 1992) (noting that confusion is relevant to nominative fair use).

342. *Kelly-Brown v. Winfrey*, 717 F.3d 295, 308 (2d Cir. 2013) (“fair use . . . requires consideration of facts outside of the complaint and thus is inappropriate to resolve on a motion to dismiss.”); see also 2 MCCARTHY, *supra* note 2, § 11:49 (“Because classic fair use is an affirmative defense, it is normally not appropriate for consideration on a . . . motion to dismiss for failure to state a claim.”).

343. Gideon Parchomovsky & Alex Stein, *Catalogs*, 115 COLUM. L. REV. 165, 178 (2015).

344. See Welkowitz, *supra* note 319, at 168 (referencing Fed. R. Evid. 301).

345. 15 U.S.C. § 1052(e)(5).

346. See *Abercrombie & Fitch Co. v. Hunting World, Inc.*, 537 F.2d 4, 9–10 (2d Cir. 1976) (explaining the limitations on generic and descriptive marks).

commentary, parody, or education should fall within safe harbors.³⁴⁷ Critiquing products or corporate behavior requires us to use them.³⁴⁸ Therefore, the first safe harbor should be expressive uses of protected trademarks.

The second safe harbor is referential uses of trademarks. Nominative fair use (referring to the trademark holder or its products) should not trigger liability.³⁴⁹ For example, rivals and repair services need to make referential uses to compete and advertise their services to the public.³⁵⁰ The law currently recognizes comparative use as a defense, but it should go further and offer a safe harbor to these uses.³⁵¹

Recognizing that the *Polaroid* factors are a “bad fit” in nominative fair use cases, one court instructed that future courts should consider fair use alongside the *Polaroid* factors when considering a claim of nominative fair use.³⁵² Likelihood of confusion is relevant to determining whether the use is objectively fair and whether defendants use the term “other[] than as a mark.”³⁵³ Likewise, nominative fair use folds confusion into determining whether an expressive use “explicitly misleads” consumers or whether the use falsely suggests a source or sponsorship.³⁵⁴

347. See Andy Greene, *Nathan Fielder Talks ‘Dumb Starbucks’ and Pranking Instagram*, ROLLING STONE (July 24, 2014), <http://www.rollingstone.com/movies/news/nathan-fielder-talks-dumb-starbucks-and-pranking-instagram-20140724> [<https://perma.cc/Q3BL-R96R>] (noting that parody laws allowed comedian Nath Fielder to “open up a near perfect replicate of a Starbucks just as long as he put the word ‘dumb’ before everything in the store, down to CDs labeled ‘Dumb Nora Jones Duets’”).

348. See *New Kids on the Block v. News Am. Publ’g, Inc.*, 971 F.2d 302, 307 (9th Cir. 1992) (“Much useful social and commercial discourse would be all but impossible if speakers were under threat of an infringement lawsuit every time they made reference to a person, company or product by using its trademark.”).

349. See *e.g.*, *Zatarains, Inc. v. Oak Grove Smokehouse, Inc.*, 698 F.2d 786, 788 (5th Cir. 1983) (rivals allowed to use “fish fry” to describe their own batter mixes even when doing so creates some likelihood of confusion with owners’ FISH-FRI trademark).

350. See *e.g.*, *Toyota Motor Sales, U.S.A., Inc. v. Tabari*, 610 F.3d 1171, 1180–82 (9th Cir. 2010) (allowing automobile broker specializing in facilitating Lexus purchases to use LEXUS mark as part of domain name).

351. See *e.g.*, *Smith v. Chanel, Inc.*, 402 F.2d 562, 563 (9th Cir. 1968) (holding that truthful comparative advertising is not trademark infringement).

352. *Int’l Info. Sys. Sec. Certification Consortium, Inc. v. Sec. Univ., LLC*, 823 F.3d 153, 168 (2d Cir. 2016); see also *Coty Inc. v. Excell Brands, LLC*, 277 F. Supp. 3d 425, 457 (S.D.N.Y. 2017) (“Acourt considering a claim of nominative fair [use] should consider three factors in addition to the standard *Polaroid* factors.”).

353. See *e.g.*, *KP Permanent Make-Up, Inc. v. Lasting Impression I, Inc.*, 543 U.S. 111, 118 (2004).

354. William McGeeveran, *Rethinking Trademark Fair Use*, 94 IOWA L. REV. 49, 95–97, 100–04 (2008).

Safe harbors like those for expressive and descriptive uses allow courts to dispose of the likelihood of confusion cases more simply and justly. For example, uses that mirror the conventional way descriptive terms are used in ordinary language give prospective users an advantage in establishing the protected use and exiting litigation early, thereby avoiding high litigation costs. In addition, they help carve out pockets of strong protection and guide the development of trademark rights in other areas such as merchandising rights, without giving owners the right to rely upon the likelihood of confusion to justify its approval. Within this framework, it is also worth considering a safe harbor beyond descriptive or expressive fair uses that provide small businesses and nonprofits like those described above with an effective and low-cost way to deflect policing by overzealous trademark owners.

III. OBSERVATIONS AND IMPLICATIONS

The final Part addresses three issues. First, it observes that courts combine factors and analyze them together. Occasionally they do so overtly. Factor folding occurs across all the factors, including the troika of mark similarity, actual confusion, and competitive proximity. A strong showing on one factor may prevent the need to show another. Second, decision-makers tend to start limiting the factors that they choose to consider when confronted with complex decision processes. At some point, decision-makers will stop analyzing new information and instead commit to a decision and then work backward to vindicate it. This adaptation has allowed 63% of litigants to receive an early resolution on the merits. Third, it explains how the empirical analysis provides a blueprint for algorithmic adjudication using AI, taking the reader from conception to execution to identifying and addressing its limitations.

A. *Factor Folding*

While likelihood of confusion factors may present themselves as discrete categories, the dataset reveals that courts do not regard them as such. Courts instead combine factors and analyze them together. This is called “folding” and in likelihood of confusion analyses the courts notoriously fold the factors together, using the presence of one factor as a proxy analysis for another.

For instance, in *J-B Weld Co., LLC v. Gorilla Glue Co.*,³⁵⁵ the Eleventh Circuit used similarity as a proxy for intent.³⁵⁶ Instead of making an adverse finding outright, a court may sometimes shift the plaintiff's burden of proof to require the defendant instead to disprove bad faith.³⁵⁷ The direction of a court's substitution bias is not a one-way street. On occasion, courts leaned on lack of evidence of actual confusion to vindicate imitation of successful product features.³⁵⁸ In this way, the likelihood of confusion factors operate not as independent elements along orthogonal lines but as a sliding scale: the more closely the products compete, the more likely it is that a new product whose design arrogates the atypical qualities of the old product will confuse consumers.

Sometimes the combination is obvious. For instance, courts treat actual confusion as an indicator of mark strength.³⁵⁹ One court explained that "[i]f buyers are confused between two sources, then this also means that they must have recognized plaintiff's designation as a trademark and associated it only with the plaintiff."³⁶⁰ Another court observed that "where the parties' marks are identical and their goods are in very close competitive proximity, a highly sophisticated consumer may be the most vulnerable to confusion."³⁶¹

This blending was not confined to the likelihood of confusion factors but extended to fair use. One court explained actual confusion gets to "the heart of the nominative fair use situation."³⁶² At other times, the logical connection is more tenuous, suggesting a negative

355. 978 F.3d 778, 790 (11th Cir. 2020).

356. *Id.* ("[W]here a defendant attempts to copy a plaintiff's product 'as closely as possible' and uses the plaintiff's product design as a model, it may be 'inferred that [defendant] purposely chose a mark which was very similar to [plaintiff's] in order to benefit from the reputation [plaintiff]'s mark had already achieved.'")

357. *Kiki Undies Corp. v. Promenade Hosiery Mills, Inc.*, 411 F.2d 1097, 1101 (2d Cir. 1969) ("[W]here the allegedly infringing mark is identical to the registered mark, and its use began subsequent to the plaintiff's trade-mark registration, the defendant must carry the burden of explanation and persuasion.")

358. *George Basch Co. v. Blue Coral, Inc.*, 968 F.2d 1532, 1541 (2d Cir. 1992) ("Absent confusion, imitation of certain successful features in another's product is not unlawful . . .").

359. *Vital Pharms., Inc. v. Monster Energy Co.*, No. 19-60809-CIV, 2021 WL 3371942, at *43 (S.D. Fla. Aug. 3, 2021) (stating that several courts have held that "actual confusion is an indicium of secondary meaning"); *Am. Sci. Chem., Inc. v. Am. Hosp. Supply*, 690 F.2d 791, 793 (9th Cir. 1982).

360. 2 MCCARTHY ON TRADEMARKS § 15:11; *see also* *Popular Bank of Fla. v. Banco Popular de P.R.*, 9 F. Supp. 2d 1347, 1358 (S.D. Fla. 1998) (making the same point).

361. *Kohler Co. v. Bold Int'l FZCO*, 422 F. Supp. 3d 681, 730 n.20 (E.D.N.Y. 2018).

362. *Juul Labs, Inc. v. 4X PODS*, 509 F. Supp. 3d 52, 64 (D.N.J. 2020).

form of coherence-based reasoning. For instance, in ascertaining mark similarity, one court considered the nature of the purchasing process, reasoning that where marks are similar but used in different contexts or on different visual displays, the risk of confusion is minimized.³⁶³

This practice of factor folding happens across all factors, including with the trio of key factors of mark similarity, actual confusion, and competitive proximity, where a strong showing on one factor may be sufficient. A strong showing on one factor may prevent the need to show another.³⁶⁴ Factors thus trump each other, with competitive proximity often trumping mark similarity without explaining why one factor should take precedence over another.³⁶⁵

Notably, only 9% of cases in the dataset expressly acknowledge “folding” factors. Most do not, regardless of the procedural posture in the case. This makes it more difficult for appellate courts and commentators to hold lower courts accountable for their analysis when this “folding” occurs. This phenomenon underscores the importance of minimizing coherence-based reasoning by having courts focus on a few factors when making likelihood of confusion determinations.

B. *Early Off-Ramps*

Courts generally agree that “application of the factors is a highly fact-intensive inquiry both as to the assessment of the evidence concerning each factor and as to the overall synthesis of factors and the evidence.”³⁶⁶ A context-specific inquiry guides courts towards the material aspects of product source or affiliation germane to the consuming public’s understanding.³⁶⁷ Given their marching orders, one might expect judges to weigh the likelihood of confusion factors

363. *Sensient Techs. Corp. v. SensoryEffects Flavor Co.*, 613 F.3d 754, 765 (8th Cir. 2010); *E.A. Sween Co. v. A & M Deli Express, Inc.*, No. 17 CV 2514, 2018 WL 1283682, at *4 (E.D.N.Y. Mar. 9, 2018).

364. *See e.g., ConAgra, Inc. v. George A. Hormel & Co.*, 990 F.2d 368, 371 (8th Cir. 1993) (“[W]hen ‘products are closely related, less similarity in trademarks is necessary to support a finding of infringement.’” (quoting *SquirtCo v. Seven-Up Co.*, 628 F.2d 1086, 1091 (8th Cir. 1980))); *Select Comfort Corp. v. Baxter*, 996 F.3d 925, 934 (8th Cir. 2021) (“[T]he relative importance of any given factor is influenced greatly by how the other factors might apply.”).

365. *Reply All Corp. v. Gimlet Media, LLC*, 843 F. App’x 392, 396 (2d Cir. 2021) (“[W]hile the two marks undoubtedly share aural and typographic similarities, they are unlikely to appear in the marketplace in a similar manner.”).

366. *Select Comfort Corp.*, 996 F.3d at 933–34.

367. *Id.* at 934 (“Common sense is inherent in the factors, and the factors, properly applied, should try to capture a holistic view of the normal experiences for any given industry, product, or service.”).

carefully.³⁶⁸ However, as seen in this study, that is not what happens in practice, nor indeed more generally.³⁶⁹

When confronting complex decision processes, decision-makers tend to limit the factors they consider.³⁷⁰ After a certain point, judges will stop analyzing new information, instead committing to their decision first and then working backwards to rationalize it. Some courts opt for a holistic weighing of the factors rather than attempting piecemeal arithmetic.³⁷¹ Others emphasize case-by-case determination, and in so doing, underscore flexibility in applying a multitude of factors.³⁷²

To resist a movant's summary judgment motion, the non-moving party must establish, through pleadings, depositions, answers to interrogatories, admissions, and affidavits in the record that there is no genuine issue of material fact.³⁷³ To put it differently, summary judgment needs to be based on undisputed material facts that show there is "only one conclusion a trier of fact could reasonably draw."³⁷⁴ The "factors require a fact-intensive inquiry not suitable for a motion to dismiss."³⁷⁵ Some appellate courts caution district courts to only grant summary judgment on the issue of likelihood of confusion "sparingly."³⁷⁶

However, expedient determinations serve the ends of justice for both sides in litigation. Summary judgments provide a quick and

368. See *Zubulake v. UBS Warburg LLC*, 217 F.R.D. 309, 322 (S.D.N.Y. 2003) (noting the tendency towards this type of application).

369. Anthony E. Chavez, *Using Legal Principles to Guide Geoengineering Deployment*, 24 N.Y.U. ENVTL. L.J. 59, 93 (2016) ("Decision makers, however, often do not apply multi-factor—or multi-principle—tests as they are intended.")

370. Beebe, *supra* note 6, at 1601.

371. *Noble v. United States*, 231 F.3d 352, 359 (7th Cir. 2000).

372. John S. Applegate, *Worst Things First: Risk, Information, and Regulatory Structure in Toxic Substances Control*, 9 YALE J. REG. 277, 302 (1992).

373. *ServPro Intell. Prop., Inc. v. Blanton*, 451 F. Supp. 3d 710, 722 (W.D. Ky. 2020).

374. *Health Net v. U.S.A. Healthnet, Inc.*, No. CV 92-3925 KN, 1993 WL 209558, at *1 (C.D. Cal. May 12, 1993).

375. *GeigTech E. Bay LLC v. Lutron Elecs. Co.*, 352 F. Supp. 3d 265, 285 (S.D.N.Y. 2018); see also *Eliya, Inc. v. Kohl's Dep't Stores*, No. 06 Civ 195, 2006 WL 2645196, at *13 n.2 (S.D.N.Y. Sept. 13, 2006) ("[A]n application of the so-called *Polaroid* factors on this motion to dismiss would be inappropriate because it would involve premature factfinding.")

376. *Fortune Dynamic, Inc. v. Victoria's Secret Stores Brand Mgmt., Inc.*, 618 F.3d 1025, 1039 (9th Cir. 2010); see also *Rearden LLC v. Rearden Com., Inc.*, 683 F.3d 1190, 1210 (9th Cir. 2012) ("Given the open-ended nature of this multi-prong inquiry, it is not surprising that summary judgment on 'likelihood of confusion' grounds is generally disfavored.")

inexpensive off-ramp for parties to dispose of a case when no real issues call for a trial. The ability of courts to wield this important judicial tool protects defendants against frivolous lawsuits and plaintiffs from incurring unnecessary costs.³⁷⁷ Streamlining the test by consolidating and trimming down the factors will enable courts to get to the heart of the inquiry expeditiously. Simplifying the likelihood of confusion lowers the temperature and makes it easier for owners to determine when to protect their interests.

Many courts are more willing to move ahead with summary judgments even where factors are in dispute and the evidence is not obvious on the basis that “as with any other issue of fact, summary judgment remains appropriate when no jury reasonably could have ruled in the non-moving party’s favor.”³⁷⁸ Others are willing to do so when most of the relevant factors weigh in the movant’s favor, including at least one “key factor.”³⁷⁹ Yet others maintain “a finding of a likelihood of confusion ‘need not be supported by a majority’ of the digits.”³⁸⁰ Non-movants resisting summary judgment must show “how additional discovery on these issues would create a genuine issue of fact” material to movants’ claim for trademark infringement.³⁸¹

In a Ninth Circuit case from the dataset, the court held as a matter of law that the trademark owner was entitled to summary judgment where the marks were identical, the goods were related, and the marketing channels overlapped.³⁸² A small set of key factors helps structure the likelihood of confusion inquiry and gives notice of pertinent issues and relevant evidence; this creates a more solid basis

377. Elaine Kussurelis, *Canada’s Summary Trial Procedure: A Viable Alternative to Summary Judgment on Trademark Likelihood of Confusion Actions in the United States*, 50 U. MIAMI INTER-AM. L. REV. 165, 168 (2019) (observing summary judgments “can be a powerful trademark litigation weapon for either plaintiffs or defendants”).

378. *RXD Media, LLC v. IP Application Dev. LLC*, 986 F.3d 361, 375 (4th Cir. 2021); see also *EST Inc. v. Royal-Grow Prods., LLC*, 526 F. Supp. 3d 943, 956–57 (D. Kan. 2021) (“The evidence is far from one-sided and leads to no obvious answer.”); *Collins v. U.S. Dep’t of Veterans Affs.*, 497 F. Supp. 3d 885, 894 (S.D. Cal. 2020) (“[S]ummary judgment is still proper in trademark infringement cases where, as here, no genuine issue of material fact exists.”).

379. See, e.g., *RXD Media, LLC*, 986 F.3d at 375 (“Based on the record before us, we hold that a jury could not have reasonably concluded that RXD’s use of the ‘ipad’ mark was unlikely to cause consumer confusion.”).

380. *Future Proof Brands, LLC v. Molson Coors Beverage Co.*, 982 F.3d 280, 298 (5th Cir. 2020).

381. *Better Angels Soc’y, Inc. v. Inst. for Am. Values, Inc.*, 419 F. Supp. 3d 765, 780 (S.D.N.Y. 2019).

382. *Nissan Motor Co. v. Nissan Comput. Corp.*, 378 F.3d 1002, 1019 (9th Cir. 2004).

for predicting case outcomes and may even be sufficient to justify an inference that there is a likelihood of confusion.³⁸³

Similarly, courts can also rule on a motion to dismiss under Rule 12(b)(6) if the complaint contains facts that state a claim to relief that is plausible on its face that the plaintiff's claims are barred as a matter of law.³⁸⁴ Where plaintiffs can substantiate the plausibility of their claims, courts will deny the motion to dismiss.³⁸⁵ Plaintiffs have successfully done so on competitive proximity,³⁸⁶ actual confusion,³⁸⁷ mark strength,³⁸⁸ mark similarity, and "bridging the gap."³⁸⁹

This Article reveals that courts rely on a small number of factors to economize their decisions to give parties an early off-ramp. With either party as the movant or cross-motion, summary judgments comprised 48% of cases in the dataset. Motions to dismiss by either party made up 10% of cases, and other postures, mostly preliminary injunction motions, made up 21%. These collectively indicate that 79% of litigants seek an early resolution. Of these, 63% received a resolution with either plaintiff or defendant winning on the merits.

Some district courts treat the likelihood of confusion as a question of fact, requiring proof of each element of each factor and categorically precluding summary judgment.³⁹⁰ Others treat the

383. *Future Proof Brands, LLC*, 982 F.3d at 298.

384. *Uber Inc. v. Uber Techs., Inc.*, 521 F. Supp. 3d 455, 462 (S.D.N.Y. 2021); *see also id.* at 464 ("At the Rule 12(b)(6) stage, the *Polaroid* analysis is limited to the facts alleged in the Complaint and any documents integral thereto.").

385. *See Ashcroft v. Iqbal*, 556 U.S. 662, 678 (2009) (holding that "[t]o survive a motion to dismiss," a complaint need only "contain sufficient matter, accepted as true, to 'state a claim to relief that is plausible on its face.'" (quoting *Bell Atlantic Corp. v. Twombly*, 550 U.S. 544, 570 (2007))).

386. *Uber Inc.*, 521 F. Supp. 3d at 464 ("It suffices to note that the Complaint plausibly alleges that defendants' services are in competitive proximity with the plaintiff's graphic design and marketing services.").

387. *Id.* ("[T]he Complaint's descriptions of confusion among businesses, official bodies and members of the public provide some factual support for the plausibility of plaintiff's claims.").

388. *Id.* ("Among the other *Polaroid* factors, the Complaint plausibly alleges that plaintiff's Uber mark is arbitrary and distinctive, that the parties' marks are similar, and that defendants' putative entry into the display-advertising market may result in plaintiff bridging the 'gap' between its services and those of defendants.").

389. *Id.*

390. *Flower Mfg., LLC v. CareCo, LLC*, 466 F. Supp. 3d 797, 820 (N.D. Ohio 2020) ("Because each of Flower's claims requires proof of that element, they cannot survive summary judgment.").

likelihood of confusion as a matter of law, whether outright³⁹¹ or paying lip service to its factual dimensions.³⁹² Appellate courts are split along the same lines, either reviewing the lower courts' "ultimate conclusion about likelihood of success for clear error,"³⁹³ or "de novo, using the same legal standards [the lower court] employed."³⁹⁴

Even among those who profess fidelity to the fact/law distinction in theory, applying that distinction in practice is not easy. As the Sixth Circuit put it, "[a]ny dispute about the evidence that pertains to the eight factors presents a factual issue . . . [and] whether a given set of foundational facts establishes a likelihood of confusion is a legal conclusion."³⁹⁵ In contrast, the Second Circuit expressed that determining whether one of the *Polaroid* factors favors one party or another is a legal judgment reviewed de novo.³⁹⁶

The data also revealed a steady increase in the affirmance rate of lower courts' decisions between 2016 and 2021. This is because appellate courts generally defer to lower court finding of facts but give less deference to questions of law.³⁹⁷ But, more significantly, appellate courts seem either unaware or complicit in the practice of lower courts folding some factors and ignoring others. It would be interesting and worth further study to see whether this is a practice of "wink-and-nod" between the lower and appellate courts or if this state of affairs was purely coincidental.

391. *FCOA, LLC v. Foremost Title & Escrow Servs., LLC*, 416 F. Supp. 3d 1381, 1387 (S.D. Fla. 2019) ("[I]n trademark infringement cases, courts in this Circuit have decided the issue of likelihood of confusion as a matter of law.").

392. *Yellowfin Yachts, Inc. v. Barker Boatworks, LLC*, 898 F.3d 1279, 1289 (11th Cir. 2018) ("Although likelihood of confusion is a question of fact, it may be decided as a matter of law.").

393. *Future Proof Brands, LLC v. Molson Coors Beverage Co.*, 982 F.3d 280, 298 (5th Cir. 2020); *AWGI, LLC v. Atlas Trucking, Co., LLC*, 998 F.3d 258, 264 (6th Cir. 2021) ("We review the district court's finding of fact for clear error . . ." (quoting *Premium Freight Mgmt., LLC v. PM Eng'g Sols., Inc.*, 906 F.3d 403, 406 (6th Cir. 2018))).

394. *Yellowfin Yachts, Inc.*, 898 F.3d at 1289.

395. *Progressive Distrib. Servs., Inc. v. United Parcel Serv., Inc.*, 856 F.3d 416, 427 (6th Cir. 2017); see also *Champions Golf Club, Inc. v. The Champions Golf Club, Inc.*, 78 F.3d 1111, 1116 (6th Cir. 1996) ("Whether there is a likelihood of confusion is a mixed question of fact and law.").

396. *Tiffany & Co. v. Costco Wholesale Corp.*, 971 F.3d 74, 86 (2d Cir. 2020) ("[I]nsofar as the determination of whether one of the *Polaroid* factors favors one party or another involves a legal judgment—which it often does—we must review that determination de novo.").

397. *Select Comfort Corp. v. Baxter*, 996 F.3d 925, 934 (8th Cir. 2021) ("[W]e review the likelihood of confusion determination as a finding of fact.").

Figure 15: Affirmance on Appeal

2016	2017	2018	2019	2020	2021
100.0%	87.5%	80.0%	84.6%	81.3%	81.8%

Scholars have long debated whether AI can replicate human legal reasoning.³⁹⁸ Edward Levi described how common law rules evolve in his classic text *An Introduction to Legal Reasoning*.³⁹⁹ Judges begin by identifying factors that have legal salience to explain case outcomes. Once those rules fail to yield sensible results, judges alter them accordingly.⁴⁰⁰ Legal reasoning rests on analogies, but this fact-specific method also makes developing coherence in the case of precedent an elusive task. For this reason, AI's ability to detect patterns in judicial opinions is of great interest to scholars.⁴⁰¹

At its heart, the likelihood of confusion inquiry seeks to ascertain the probability that a defendant's use of its trademarks will confuse consumers.⁴⁰² Making the likelihood of confusion more rule-like, both through the doctrinal reformation of the standard and through the application of AI, makes it easier for appeals courts to scrutinize and overturn deviant lower court decisions and allows lower courts to

398. See e.g., Stephen M. McJohn, *Review of Artificial Legal Intelligence*, 12 HARV. J.L. & TECH. 241, 243–44 (1998) (noting that the way neural networks learn through adjustment makes the technology capable of performing legal reasoning; however, the technology would not understand the reasons behind its conclusion, making it less valuable to the legal field).

399. Edward H. Levi, *An Introduction to Legal Reasoning*, 15 U. CHI. L. REV. 501, 501–03 (1948).

400. *Id.*; see also RONALD DWORKIN, *LAW'S EMPIRE* 400–03 (1986) (explaining the factors that a judge might weigh when considering whether to change the law prospectively, noting that they are bound by precedent and the integrity of law).

401. See, e.g., McJohn, *supra* note 398, at 241 (offering commentary on another scholar's contribution to discussions around the topic).

402. Leah Chan Grinvald, *Shaming Trademark Bullies*, 2011 WIS. L. REV. 625, 636 (2011) (“This liability standard refers to the probability (not the actuality or possibility) that consumers will be confused by the same or similar trademarks.”).

distinguish dubious precedent based on facts.⁴⁰³ Trademark’s troika of actual confusion, mark similarity, and competitive proximity paves the road for AI to fill the final piece of the equation to simplify the likelihood of confusion.

C. Deploying Artificial Intelligence

Legal scholarship on AI and trademarks is scarce.⁴⁰⁴ It is surprising given AI’s centrality in both the consumer marketing literature and trademark’s centrality in IP protection.⁴⁰⁵ Sonia Katyal and Aniket Kesari argued “as a general matter, that AI should be of interest to anyone studying trademarks and the role that they play in economic decision making.”⁴⁰⁶ They point to AI deployment by the government in trademark image recognition, classifying goods and services, and identifying descriptive terms.⁴⁰⁷ This Article explains how a dataset such as the one used here might be a rudimentary prototype for a grander form of AI-enabled likelihood of confusion analysis that courts and litigants might deploy in the future.

403. Frederick Schauer, *Formalism*, 97 YALE L.J. 509, 541–42 (1988) (noting errors are more easily detectable under rules).

404. Sonia K. Katyal & Aniket Kesari, *Trademark Search, Artificial Intelligence, and the Role of the Private Sector*, 35 BERKELEY TECH. L.J. 501, 504 (2020) (“[S]urprisingly, very little legal scholarship has addressed the potential role for AI in the context of trademarks.”). The few examples available include Dev Gangjee, *Eye, Robot: Artificial Intelligence and Trade Mark Registers*, in TRANSITION AND COHERENCE IN INTELLECTUAL PROPERTY LAW (Niklas Bruun, Graeme B. Dinwoodie, Marianne Levin & Ansgar Ohly eds., 2021), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3467627; Anke Moerland & Conrado Freitas, *Artificial Intelligence and Trade Mark Assessment*, in ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES 49–50 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3683807.

405. See *Trademarks, Copyright and Patents: Should Business Owners Really Care About IP?*, VARNUM (May 1, 2019), <https://www.varnumlaw.com/newsroom-publications-trademarks-copyrights-and-patents-why-business-owners-should-care-about-ip> [<https://perma.cc/VST8-XF56>] (“A trademark is one of the most important business assets that a company will ever own because it identifies and distinguishes the company and its products/services in the marketplace from its competitors.”).

406. Katyal & Kesari, *supra* note 404, at 505 (“AI will fundamentally transform the trademark ecosystem, and the law will need to evolve as a result.”).

407. See, e.g., DAVID FREEMAN ENGSTROM ET AL., GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES 49–50 (2020), <https://www-cdn.law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf> [<https://perma.cc/SR9J-UEUN>] (describing a USPTO prototyping using deep learning model using an unsupervised approach to generate visually similar images from a database).

1. *Conception*

AI gives courts the capability to scour case reports to assess how past courts weighed effects and stress-test theories of confusion against real-world data.⁴⁰⁸ AI can match the results against depositions and other preprocessed evidence to provide quicker and more consistent analyses, unlike the binarily coded factors in this study.⁴⁰⁹ Principal component analysis can identify factors carrying the greatest weight in functions and zero in on the most important dimensions of datasets to show the stampeding likelihood of confusion factors.⁴¹⁰

AI expands the scope of cases so that courts can dispense cases summarily. It can significantly reduce the time and effort needed to analyze a case, and courts can apply consistently evolving legal principles, even when the facts are idiosyncratic.⁴¹¹ It can also avoid the risk of judges engaging in side-by-side mark comparison and ensure they apply the real-world purchasing context. The results from AI recommendations challenge judges' prior assumptions, providing a check against coherence-based reasoning. Simon's research shows that confronting people with merits of the opposite side reduced the effect of coherence shifts by about 50%.⁴¹² In particular, his study moderated jury instruction by expressly requesting jury members to "take some time to seriously consider the possibility that the opposite side has a better case."⁴¹³ Other legal studies similarly showed that asking lawyers to consider the weaknesses in their side or reasons that the judge might rule against them mitigated bias.⁴¹⁴

The beauty of AI-enabled likelihood of confusion analysis is that it can reach outcomes we cannot define in advance of the AI being run as "good" or "better" than the untrained neural network interrogates itself via the process of trial and error. In addition, convolutional neural networks can abstract local features from examples, for instance, by recognizing specific facts in opinions. They would also account for interactions among indicators that escape even expert

408. See Daryl Lim, *Confusion, Simplified*, BERKELEY TECH. L.J. (Forthcoming, 2022) [hereinafter Lim, *Confusion, Simplified*].

409. *Id.*

410. *Id.*

411. Pamela Samuelson, *Unbundling Fair Uses*, 77 *FORDHAM L. REV.* 2537, 2541–42 (2009).

412. Simon, *Third View of the Black Box*, *supra* note 225, at 544 (noting that "[m]ore studies are required to gain a better sense of the effects of the debiasing intervention").

413. *Id.* at 571.

414. See Linda Babcock et al., *Creating Convergence: Debiasing Biased Litigants*, 22 *L. & SOC. INQUIRY* 913, 920–21 (1997).

witnesses and contextualize and associate information with known factors to provide predictions based on untrained parameters.⁴¹⁵ Finally, unsupervised data mining algorithms can zero in on data clusters and probe those clusters to find other abstractions.⁴¹⁶

Moreover, programming the AI to maximize reward in a predetermined environment allows it to directly optimize policy performance rather than learning from old data⁴¹⁷ by updating the agent's policy using good estimates of a particular policy's advantage relative to another policy.⁴¹⁸ Conceivably, variations of the algorithm will predict litigation risk and the business implications of marketing and sales decisions.⁴¹⁹ As Dev Gangjee put it, "it is extremely tempting to be guided by clearly defined percentages of similarity."⁴²⁰ A. S. Li, A.J.C. Trappey, and C.V. Trappey sketched out how that model might work. The data set combines trademark litigation ontology and text mining to extract features from cases to build a machine-readable database like case content analysis.⁴²¹

2. Execution

Like many AI datasets, case content analysis treats the content of opinions as generic data.⁴²² Coding and counting cases imply that

415. Similarly, AI-based support vector machines (SVMs) can find relationships between sets of trademark infringement cases while handling outlier or mislabeled cases, allowing SVM to crunch abrogated case law. *See e.g.*, AURÉLIEN GÉRON, *HANDS-ON MACHINE LEARNING WITH SCIKIT-LEARN, KERAS, AND TENSORFLOW* 155–67 (Nicole Tache ed., 2d ed. 2019) (explaining how SVMs work and how they can be helpful).

416. PEDRO DOMINGOS, *THE MASTER ALGORITHM* 210 (2015).

417. *Proximal Policy Optimization, OpenAI Spinning Up*, OPENAI (2018) <https://spinningup.openai.com/en/latest/algorithms/ppo.html> [<https://perma.cc/GB72-ZWGX>].

418. *See* Brian S. Haney, *AI Patents: A Data Driven Approach* 19 CHI-KENT J. INTELL. PROP. 407, 439 (2020) (explaining the advantage function).

419. Katyal & Kesari, *supra* note 404, at 533 ("Indeed, predictive analytics can prove to be transformative in helping businesses both create and sustain a strong presence in the marketplace, predicting the outcome of filing suit, sending a cease-and-desist, articulating various claims, or deciding whether and for how much to settle. And this is just the tip of the iceberg. Imagine every aspect of a trademark claim—its probable outcome automated, calculated, predicted and ready for real-time decision-making.").

420. *See* Gangjee, *supra* note 404, at 13.

421. A. S. Li, A. J. C. Trappey, & C. V. Trappey, *Intelligent Identification of Trademark Case Precedents Using Semantic Ontology*, in *TRANSDISCIPLINARY ENGINEERING FOR COMPLEX SOCIO-TECHNICAL SYSTEMS—REAL-LIFE APPLICATIONS* (Jerzy Pokojski et al. eds. 2020).

422. Hall & Wright, *supra* note 29, at 83.

information in one opinion is potentially relevant to another.⁴²³ In a pre-AI world, an army of legal scholars might attempt to map all likelihood of confusion cases comprehensively. Indeed, they would need to endure many hours of time-consuming and repetitious reading and extracting of the necessary information to code each case, draw interferences, and report on trends, as was done in this Article's writing.

Case content analysis is suitable for automating because the same set of information must be keyed into many cells in the same case. This requires coders to eyeball each cell for accuracy given the tedious, repetitive data entry, resulting in avoidable human errors and copy-paste tasks.⁴²⁴ Nevertheless, as seen in this Article, the result is useful, capable of determining the weight courts have placed on various legal and non-legal factors, identifying which factors judges use to "stampede" others, revealing trends across time, and other relevant parametric factors that may typically escape conventional wisdom.⁴²⁵

The algorithm would pick out keywords and assign appropriate weights to each variable with AI. For example, factors "in favor of"⁴²⁶ or "favors"⁴²⁷ would signify a positive correlation to one side. Similarly, phrases like "marks are strong,"⁴²⁸ "high degree of care,"⁴²⁹ "marks are . . . identical,"⁴³⁰ and "weighs heavily in favor of,"⁴³¹ would be assigned greater, or in the case of "neutral,"⁴³² "weighs neither for

423. See Peter J. Hammer & William M. Sage, *Antitrust, Health Care Quality, and the Courts*, 102 COLUM. L. REV. 545, 561 (2002) (explaining that the coding exercise does not determine the law, but instead treats opinions as data).

424. Serena Lim & Nandini Nayar Sharma, *Document Drafting—Less Is More*, SING. L. GAZETTE (Nov. 21, 2021), <https://lawgazette.com.sg/practice/tech-talk/document-drafting-less-is-more>.

425. See e.g., Barton Beebe, *An Empirical Study of U.S. Copyright Fair Use Opinions, 1978–2005*, 156 U. PA. L. REV. 549, 591 (2008) (describing a study of cases regarding the fair use reasoning used by different courts).

426. *Fletcher's Original State Fair Corny Dogs, LLC v. Fletcher-Warner Holdings LLC*, 434 F. Supp. 3d 473, 485 (E.D. Tex. 2020).

427. E.g., *Fleet Feet, Inc. v. Nike Inc.*, 419 F. Supp. 3d 919, 943 (M.D.N.C. 2019) ("This factor favors Fleet Feet."), *appeal dismissed*, 986 F.3d 458 (4th Cir. 2021).

428. *New Balance Athletics, Inc. v. USA New Bunren Int'l Co.*, 424 F. Supp. 3d 334, 347 (D. Del. 2019).

429. *Id.* at 348.

430. *Id.* at 347.

431. *Better Angels Soc'y, Inc. v. Inst. For Am. Values, Inc.*, 419 F. Supp. 3d 765, 777 (S.D.N.Y. 2019); *Life After Hate, Inc. v. Free Radicals Project, Inc.*, 410 F. Supp. 3d 891, 908 (N.D. Ill. 2019).

432. *Illinois Tool Works Inc. v. J-B Weld Co.*, 419 F. Supp. 3d 382, 399 (D. Conn. 2019).

[n]or against,”⁴³³ “slightly in favor of,”⁴³⁴ would be afforded less weight. The algorithm would also recognize and capture variables like rivalry (“direct competitors”⁴³⁵).

Automation saves dataset preparers a substantial amount of time. Studies on automating conveyancing work show a time savings of 90%.⁴³⁶ Less skilled and lower-cost staff can quickly and accurately generate datasets, lowering costs, time, and effort to produce complex datasets, freeing up scholars to focus on higher-value work.⁴³⁷ The user selects a smart template and answers a questionnaire presented by the template to generate an opinion.⁴³⁸ The AI then uses the training data to assemble a custom opinion.⁴³⁹

Likelihood of confusion opinions contain logic-dependent conditional clause variations which incorporate the factors. The algorithm could compare the qualitative and quantitative factors presented in each case to its markers as a first step. Cases presenting the same set of facts would reach the same outcome as precedential cases presenting the same set of markers. AI will need to specify the weight of factors not expressly entailed by rules or precedents. Once algorithms produce their recommendation, judges could accept or reject the AI’s recommendation, like how Amazon consumers choose to make another purchase based on Amazon’s recommendations of their earlier purchases and browsing history.⁴⁴⁰

The algorithm randomly plays out certain results, learns—with input from data scientists in each iteration—adjusts its weights and parameters, and chooses advantageous moves with increasing finesse.⁴⁴¹ The feedback loop causes the algorithm’s nodes to change their weights, so case precedents refined by new case law and market

433. *New Balance Athletics, Inc.*, 424 F. Supp. 3d, at 349.

434. *Delta Forensic Eng’r, Inc. v. Delta V Biomechanics, Inc.*, 402 F. Supp. 3d 902, 910 (C.D. Cal. 2019).

435. *Smartling, Inc. v. Skawa Innovation Ltd.*, 358 F. Supp. 3d 124, 139 (D. Mass. 2019).

436. *Lim & Sharma*, *supra* note 424.

437. *Id.*

438. *Id.*

439. *Id.*

440. See SPANDANA SINGH, *NEW AMERICA, WHY AM I SEEING THIS? HOW VIDEO AND E-COMMERCE PLATFORMS USE RECOMMENDATION SYSTEMS TO SHAPE USER EXPERIENCE* 22 (Mar. 2020) (discussing how Amazon’s recommendation system drives user purchases and contributes to revenue generation on the platform).

441. See TERRENCE J. SEJNOWSKI, *THE DEEP LEARNING REVOLUTION* 20 (2018) (describing how the AlphaZero used machine learning algorithm to play chess).

data may eventually yield a different, better outcome over time.⁴⁴² This allows adjudicating to become less a question of “ideology plus facts plus law equal the outcome” and more a question of whether the data supports the parties’ legal outcome or if on appeal, one that the lower court advanced.

3. *Limitations*

As with any AI system, there are limitations, some generic, some specific, that its implementers need to keep in mind. First, this Section identifies the main limitations pertinent to the discussion: “garbage-in, garbage-out,” biases, contextualizing purchasing conditions, and coding challenges. Then, this Section discusses each one in turn.

a. *“Garbage-in, garbage-out”*

First, the saying “garbage-in, garbage-out” applies to the training dataset. The algorithm applies the judge’s expertise through the opinions coded in the training data while minimizing unreliability.⁴⁴³ The case law may be doctrinally flawed but remain good law. Nonetheless, the algorithm can implement the likelihood of confusion factors more consistently than both the human judges who decide the precedential cases in the dataset and the judges applying those precedents. Moreover, judges adjudicating live cases can compare the model’s prediction with the ground truth and adjust the model’s parameters, minimizing the error between these two values over time. As algorithms gain additional knowledge about the probabilities of occurrence, ambiguity disappears, and the choices become clearer.⁴⁴⁴

Scholars and AI service providers agree that AI augments human decision-making and does not displace it.⁴⁴⁵ As LawPanel’s founder put

442. See Stephen McJohn & Ian McJohn, *Fair Use and Machine Learning*, 12 NE. U. L. REV. 99, 135 (2020).

443. Dawes, *supra* note 216, at 575 (“[A] linear model distills underlying policy . . . from otherwise variable behavior (e.g., judgments affected by context effects or extraneous variables).”).

444. See Gary Charness & Dan Levin, *When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect*, 95 AM. ECON. REV. 1300, 1300 (2005) (describing a comparable heuristic form of processing new information).

445. See Gangjee, *supra* note 404, at 11 (“Experience till date therefore suggests that AI algorithms are intended to augment human judgment—to effectively sift through ever increasing volumes of registration data—and not to replace it.”); see also COMPUMARK, ARTIFICIAL INTELLIGENCE, HUMAN EXPERTISE: HOW TECHNOLOGY AND TRADEMARK EXPERTS WORK TOGETHER TO MEET TODAY’S IP CHALLENGES 5 (2018) (“While AI and neural networks will play an expanding role in CompuMark solutions . . . they are intended to complement, not replace, human analysts.”).

it, “AI will speed up legal research, but it will not replace advice formulation . . . [since it] only works on repetitive tasks in a very tightly-defined domain.”⁴⁴⁶ Nevertheless, Katyal and Kesari are optimistic that the gap can be closed as data scientists enrich the dataset with more data points and human-AI teams.⁴⁴⁷ They report how experts are continuing to highlight the need for human oversight and participation, particularly when it comes to complex cognitive tasks in trademark doctrines.⁴⁴⁸

On appeal, the variability of decisions can reveal some idea of the extent of noise. A three-judge circuit appeals court or nine-justice Supreme Court bench would provide an additional check. Salib observes that “there will be an adjustment period as courts develop doctrine about what constitutes credible scientific practice in algorithmic design. Such bumps on the road, however, are the cost of admission if generalist judges are to continue playing any major role governing our increasingly complex world.”⁴⁴⁹

b. Biases

Second, system architects need to address data biases in adopting the technology and in deploying AI. For example, with supervised machine learning, humans classify the data. This introduces bias, such as training an AI on the similarity of signs. One trainer might determine a similarity between two given signs, while another might not.⁴⁵⁰ As a result, AI may replicate and perpetuate data biases.⁴⁵¹

Coding is not value-neutral, and biases may seep into the algorithmic code, filtering into training data and the weights judges assign to the algorithm.⁴⁵² Bias could also come from the algorithms being trained

446. Tim Lince, “No Imminent AI Apocalypse”—Tech Expert Rejects Predictions of Mass Job Losses in Trademark Industry, *WORLD TRADEMARK REV.* (Feb. 1, 2018), <https://www.worldtrademarkreview.com/brand-management/no-imminent-ai-apocalypse-tech-expert-rejects-predictions-mass-job-losses> [<https://perma.cc/RAF4-FHPV>].

447. See Katyal & Kesari, *supra* note 404, at 526.

448. *Id.* at 533.

449. Peter N. Salib, *Artificially Intelligent Class Actions*, 100 *TEX. L. REV.* 519, 552 (2022).

450. See Lim, *Confusion, Simplified*, *supra* note 408.

451. See *id.*

452. See Dan L. Burk, *Algorithmic Fair Use*, 86 *U. CHI. L. REV.* 283, 283 (2019) (describing how design values of algorithms can reflect biases); see, e.g., David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 *U.C. DAVIS L. REV.* 653, 669–701 (2017) (explaining that because humans make algorithms and humans have bias, the biases of humans are reflected in both algorithms themselves and how humans use them).

using biased data, such as prior decisions from judges who are biased themselves, and from the way humans interpret the data produced by AI systems.⁴⁵³ In addition, reinforcement learning techniques may embed bias, raising the risk of what Thomas Nachbar labeled “snowballing unfairness.”⁴⁵⁴ Codes are based on earlier program decisions and the constant integration of new information, prompts a continual search for purpose.⁴⁵⁵

Moerland and Freitas provide an example of bias in action:

[W]hen teaching an AI to establish a pattern of similarity of signs, one could easily ascertain a similarity between two given signs, while someone else would not. Even if case law regarding similarity of signs is used as training data, courts sometimes come to differing outcomes for the same cases. Bias in data will be replicated when used by the AI technology, as it lacks the ability to filter out slightly incorrect interpretations.⁴⁵⁶

The lack of a standardized method to weigh factors systematically exacerbates the risk of bias. The likelihood of confusion factors have no weights assigned, eroding the ability to apply the tests objectively or in a manner that can be replicated.⁴⁵⁷ AI helps integrate data and provides a statistical prediction based on input variables. Humans are superior at selecting and coding information but poor at integrating it.⁴⁵⁸

Daniel Kahneman, Cass Sunstein, and Olivier Sibony recommend assigning probabilities rather than absolute values or binary “yes” or “no” judgments.⁴⁵⁹ Numerical thresholds could help adjudicate infringement cases. For example, computer scientists could build a model that requires judges to rate the three core likelihood of confusion factors on a scale of 0–10. If the marks were completely different, the judge would rate it ‘0’ (the lowest rating possible), but if

453. See Gangjee, *supra* note 404, at 11 (“[W]here the data for a machine learning approach is derived from judicial content analysis—past decisions by human tribunals where factors are coded and correlations derived—the algorithm will behave like the human decision maker it is modelled after, warts and all.”).

454. Thomas B. Nachbar, *Algorithmic Fairness, Algorithmic Discrimination*, 48 FLA. STATE L. REV. 509, 522 (2021).

455. *Id.* at 548.

456. Moerland & Freitas, *supra* note 404, at 282.

457. See *Menard, Inc. v. Comm’r of Internal Revenue*, 560 F.3d 620, 622–23 (7th Cir. 2009) (“Multifactor tests with no weight assigned to any factor are bad enough from the standpoint of providing an objective basis for a judicial decision; multifactor tests when none of the factors is concrete are worse.”).

458. See, e.g., Dawes, *supra* note 216, at 573.

459. See KAHNEMAN, SIBONY & SUNSTEIN, *supra* note 176, at 218.

the marks were simple counterfeits, the judge would rate it '10' (the highest rating). Thus, the algorithm would set a numerical threshold for finding confusion that maps to case law and the balance of probabilities. Over time, the algorithm would provide more granular information about the characteristics driving outcomes in likelihood of confusion cases. In this way, the algorithm would imitate judges, granting a low score to a particular factor and a consequently lower success rate to plaintiffs.

Daniel Kahneman, Cass Sunstein, and Olivier Sibony also recommend relying more heavily on rules like judicial sentencing guidelines.⁴⁶⁰ The trio of factors again provides that framework. Importantly, the results from AI recommendations challenge judges' prior assumptions, providing a check against coherence-based reasoning.⁴⁶¹ For instance, confronting people with merits of the opposite side reduced the effect of coherence shifts by about 50%.⁴⁶² Legal studies similarly showed that asking lawyers to consider the weaknesses in their side or reasons that the judge might rule against them mitigated bias.⁴⁶³

Finally, to address the issue of "snowballing unfairness," flooding the system with voluminous data may help. As Moerland and Freitas note, "with large amounts of data, incidental bias may not influence the rule that the AI learns from the data."⁴⁶⁴ They reassuringly report that "[AI] training is continuous and subject to high standards of reliability. Error measures are used as well as pilot studies on unseen data to determine how the AI tool performs its tasks."⁴⁶⁵

Done well, trademark algocracy will minimize biases from human decision-making without compounding those biases with its own.⁴⁶⁶ In the years ahead, ethics teams will likely become an essential department in antitrust agencies and economic consultancies such as finance, legal, marketing, and human resource departments. These

460. *Id.* at 258.

461. *See generally* Lim, *Saving Substantial Similarity*, *supra* note 281, at 634 (discussing the safety valve role of juries in challenging the assumptions of judges).

462. Simon, *Third View of the Black Box*, *supra* note 225, at 544 (noting that "[m]ore studies are required to gain a better sense of the effects of the debiasing intervention").

463. Babcock et al., *supra* note 414, at 920–21 (describing results from an experiment which showed that when subjects consciously considered weaknesses in their cases, their biases were mitigated).

464. *See* Moerland & Freitas, *supra* note 404, at 282.

465. *Id.* at 281.

466. *See* Meghan J. Ryan, *Secret Conviction Programs*, 77 WASH. & LEE L. REV. 269, 286 (2020) (describing how computer programs can effectively help judges avoid injecting their own biases when making judicial decisions).

teams can help decision-makers weigh the benefits and harms of AI procedures and recommendations, flag their implications, develop guidelines, and help clarify ethical conflicts.⁴⁶⁷

c. Contextualizing the purchasing conditions

Third, the algorithm needs to replicate how a human perceives a mark in the marketplace.⁴⁶⁸ As we saw in Part I, a human fact-finder trying to contextualize the marketplace faces a difficult task. AI has the additional burden of delivering on that promise of an objective assessment.⁴⁶⁹ The algorithm will need to account for the relevant consumer type, competitors, circumstances of purchase, and the end-use of the product. The AI tool is unlikely to contextualize and juxtapose them against case law to compare situations and determine a likely outcome. The data scientist will need to acquire this information and structure it in a manner that the algorithm can automatically process.

One way to contextualize purchasing conditions is for AI to maximize a preset reward without the need for continual human supervision. In this case, the reward is whether, on balance, consumers would be confused.⁴⁷⁰ Instead, the algorithm chooses an action in the environment's initial state—representing a moment in time—randomly explores the environment, gathers information about the environment, develops an optimal policy, and optimizes performance by “expressing the relationship between the value of a state and the values of future states.”⁴⁷¹


467. Lim Sun & Jeffrey Chan Kok Hui, *Moving AI Ethics Beyond Guidelines*, STRAITS TIMES (Dec. 16, 2020), <https://www.straitstimes.com/opinion/moving-ai-ethics-beyond-guidelines-0> [<https://perma.cc/R4WD-76TD>].

468. Moerland & Freita, *supra* note 404, at 284 (“This leads us to the finding that the assessment is one of degree and requires reasoning from the perspective of the relevant public. It is questionable as to how far AI technology can reflect this human-centric approach.”).

469. Katyal & Kesari, *supra* note 404, at 532 (“Others have expressed similar concerns, noting that determining trademark distinctiveness, the relevant public, the proper classification of goods and services, among other elements, are so subjective that they pose challenges to the development of AI in trademark law.”).

470. See Haney, *supra* note 418, at 430 (describing how reward can act as a feedback mechanism).

471. *Id.* at 437.

Figure 16: Reinforcement Learning⁴⁷²

As it continues to the next state, the agent receives a reward and a set of choices, the algorithm selects an action, and the environment returns a reward and the next state.⁴⁷³ The reward teaches the algorithm what to do and formalizes the goal's idea.⁴⁷⁴ Through this iteration, it learns to take actions optimizing a reward, which would be, say, mark similarity.⁴⁷⁵ In essence, the total reward mirrors the legal “algorithm” we call trademark law’s likelihood of confusion.

This feature allows the algorithm to navigate dynamic market environments without stopping the environment before computing.⁴⁷⁶ To the extent variables in its dataset need modification, AI training techniques use autoencoders to update word embeddings, machine translation, document clustering, sentiment analysis, and paraphrase detection.⁴⁷⁷ Stacking autoencoders on top of each other allows the first autoencoder to focus on encoding features at one level of abstraction.⁴⁷⁸ The next autoencoder uses the earlier output to recognize fact patterns and encode more abstract features.⁴⁷⁹ Defining features broadly helps avoid overfitting, which happens when the

472. RICHARD S. SUTTON & ANDREW G. BARTO, REINFORCEMENT LEARNING: AN INTRODUCTION 48 (2018).

473. EUGENE CHARNAK, INTRODUCTION TO DEEP LEARNING 113 (2018).

474. *Id.*

475. Jennifer Barry, Daniel T. Barry, & Scott Aaronson, *Quantum Partially Observable Markov Decision Processes*, PHYS. REV., No. 032311, 2014, at 1, 2, <https://journals.aps.org/prx/abstract/10.1103/PhysRevA.90.032311>.

476. *Id.*

477. See Venkata Krishna Jonnalagadda, *Sparse, Stacked and Variational Autoencoder*, MEDIUM (Dec. 6, 2018), <https://medium.com/@venkatakrishna.jonnalagadda/sparse-stacked-and-variational-autoencoder-efe5bfc73b64> [<https://perma.cc/ZH2D-JTNW>] (describing the various types of and uses for autoencoders).

478. *Id.*

479. Gideon Lewis-Kraus, *The Great A.I. Awakening*, N.Y. TIMES, (Dec. 14, 2016), <https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html> [<https://perma.cc/46B3-8DKV>].

learner fits the function to the data.⁴⁸⁰ Overfitting also happens in legal reasoning when one ties a rule to the facts. The solution is to include more training examples and test the function against other test examples.⁴⁸¹

d. Coding challenges

Fourth, the likelihood of confusion factors do not currently lend themselves to easy coding by a machine, given the coherence-based reasoning and non-uniformity in how courts operationalize those factors, as discussed in this Article. Finally, the Eighth Circuit reminds us that “factors do not operate in a mathematically precise formula.”⁴⁸²

Again, the issue is real but not insurmountable. Courts can do their part by employing more rule-like formulations when applying the likelihood of confusion standard, such as the “rules of thumb” advanced in this Article. Courts can also standardize their lexicon, enabling them to present their judicial opinions in a way more amenable to machine learning. Finally, courts could and should also standardize their treatment of absent factors in the likelihood of confusion inquiry—do these factors favor either party and if so, in what way?⁴⁸³ This templating exercise helps rationalize and consolidate disparate variations into a reusable asset that captures and preserves the substantial knowledge of experienced judges.

Finally, the algorithm will need to distinguish between cases from courts at different levels of the judicial hierarchy. Stare decisis tells us that Supreme Court cases take precedence over court of appeal cases, which in turn take precedence over district court cases. However, empirical legal studies routinely ignore the weight stare decisis endows in coding datasets.⁴⁸⁴ It matters little if the Supreme Court or a district court looked at likelihood of confusion if the variable of interest is competitive proximity. The algorithm will need to consider judicial hierarchy, the appellate jurisdiction of regional circuit courts, and similar factors as appropriate.

480. GÉRON, *supra* note 415, at 26–28.

481. *Id.* at 29.

482. *Frosty Treats, Inc. v. Sony Comput. Ent. Am., Inc.*, 426 F.3d 1001, 1008 (8th Cir. 2005).

483. *Easy Spirit, LLC v. Skechers U.S.A., Inc.*, 515 F. Supp. 3d 47, 74 (S.D.N.Y. 2021) (explaining that absence of actual confusion favors junior user).

484. See *e.g.*, Lim, *Saving Substantial Similarity*, *supra* note 281.

CONCLUSION

Congress built a degree of indeterminacy into the likelihood of confusion standard as a feature and not a bug. Over the years, however, the jurisprudential roots of trademark law became unruly and tangled. Unwanted variability and bias in judgments cause serious problems by including complex and irrelevant factors, including financial loss and rampant unfairness. Meanwhile, simple rules and algorithms have developed with technological strides presenting big advantages over human judges. Three core factors, combined with two safe harbors and AI, would enable courts to reach consistent and accurate results. A simplified framework promotes fair play, safeguards expressive uses, and enhances access to justice.

This Article presented a contemporary empirical analysis of the likelihood of confusion factors and how they interact. Conventional wisdom teaches us that courts should comprehensively traverse each factor and that likelihood of confusion cases generally require jury determination. The data reveals that neither is true. Instead, courts provide early off-ramps to litigants by “economizing” using a handful of factors or by “folding” factors within each other. The findings also indicate which forums are pro-defendant and which are pro-plaintiff, the impact of rivalry and fair use on outcomes, and the Ninth Circuit’s dominance. This Article also showed how AI systems could use empirical studies as training data to help stakeholders make likelihood of confusion analyses. A familiar yet concise, precise, and efficient framework helps preempt, counsel, and adjudicate disputes. In this way, the likelihood of confusion standard can attain the amphibious benefits of becoming more rule-like while retaining its suppleness.



Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models

Matthew Dahl¹, Varun Magesh[†], Mirac Suzgun[‡], and Daniel E. Ho[§]

Abstract

Do large language models (LLMs) know the law? LLMs are increasingly being used to augment legal practice, education, and research, yet their revolutionary potential is threatened by the presence of “hallucinations”—textual output that is not consistent with legal facts. We present the first systematic evidence of these hallucinations in public-facing LLMs, documenting trends across jurisdictions, courts, time periods, and cases. Using OpenAI’s ChatGPT 4 and other public models, we show that LLMs hallucinate at least 58% of the time, struggle to predict their own hallucinations, and often uncritically accept users’ incorrect legal assumptions. We conclude by cautioning against the rapid and unsupervised integration of popular LLMs into legal tasks, and we develop a typology of legal hallucinations to guide future research in this area.

1 INTRODUCTION

How well do large language models (LLMs) know American case law? Modern LLMs such as OpenAI’s ChatGPT—tools trained on vast amounts of textual data to predict the next token in a sequence—are driving a transformation in the legal world, from legal education (Choi and Schwarcz 2024), to legal research (Livermore, Herron, and Rockmore 2024), to legal practice itself (Rodgers, Armour, and Sako 2023). Indeed, recent versions of these artificial intelligence (AI) models seem to excel at law-related tasks, such as first-year law school exams (Choi et al. 2022), the uniform bar exam (Katz et al. 2023), statutory reasoning (Blair-Stanek, Holzenberger, and Van Durme 2023), and issue-rule-application-conclusion (IRAC) analysis (Guha et al. 2023). But despite the revolutionary potential of these models, a key challenge remains: the issue of “hallucinations.” LLMs are liable to generate language that is inconsistent with current legal doctrine and case law, and, in the legal field, where adherence to authorities is paramount, unfaithful or imprecise interpretations of the law can lead to nonsensical—or worse, harmful and inaccurate—legal advice or decisions.

In this work, we present the first evidence documenting the nature, frequency, and correlates of these hallucinations. In doing so, we shed systematic, empirical light on a phenomenon that has so far only received anecdotal treatment in the literature. For example, much media attention has been directed toward a Manhattan lawyer who faced sanctions for using ChatGPT to generate fictional

¹ J.D. Candidate, Yale Law School, and Ph.D. Student, Yale University Department of Political Science, USA. E-mail: matthew.dahl@yale.edu.

[†] Research Fellow, Regulation, Evaluation, and Governance Lab, Stanford University, USA.

[‡] J.D. Candidate, Stanford Law School, and Ph.D. Candidate, Stanford University Department of Computer Science, USA.

[§] William Benjamin Scott and Luna M. Scott Professor of Law, Professor of Political Science, Senior Fellow at Stanford Institute for Economic Policy Research, and Faculty Director of the Regulation, Evaluation, and Governance Lab, Stanford University, USA. We thank Neel Guha, Sandy Handan-Nader, Adam T. Kalai, Peter Maldonado, Chris Manning, Joel Niklaus, Kit Rodolfa, Faiz Surani, Andrea Vallebuono, and three anonymous reviewers for helpful discussions and feedback. Our replication materials are available in the Harvard Dataverse (<https://doi.org/10.7910/DVN/V4ON8H>) and our data for future benchmarking is available on HuggingFace (https://huggingface.co/datasets/reglab/legal_hallucinations).

© The Author(s) 2024. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

© 2024 by Matthew Dahl, Varun Magesh, Mirac Suzgun, and Daniel E. Ho. Originally published in 16 J. Legal Analysis 64 (2024). Reprinted with kind permission of the authors.

case citations for a brief (Weiser 2023), or another instance where ChatGPT produced a supposed dissent authored by Justice Ruth Bader Ginsburg in the landmark gay rights case *Obergefell v. Hodges* (Romoser 2023). Even Chief Justice John Roberts, the chief justice of the U.S. Supreme Court, has weighed in on the problem, highlighting hallucinations in his 2023 report on the state of the federal judiciary and arguing that, as yet, “machines cannot fully replace key actors in court” (Roberts 2023, 6).

These impressionistic accounts, however, leave unanswered the deeper questions that legal scholars must confront as LLMs continue to grow in popularity. How much legal knowledge is actually embedded in an off-the-shelf LLM? Are LLMs equally familiar with different dimensions of the American common law system—where legal doctrine varies across courts, jurisdictions, and over time—or do they tend to hallucinate more in certain areas than others? When LLMs do hallucinate, do they disproportionately produce false information favoring certain judges or cases? And besides hallucination itself, are there other features of LLMs that legal scholars need to consider—other latent biases or behavioral tendencies that threaten to spill over into downstream applications of these models? Our study seeks to answer these questions, providing insights that are essential for evaluating LLMs’ effectiveness in general legal settings.

This research contributes to several literatures. First, there has recently been an explosion of interest in the intersection of law and technology, with a particular focus on the emergence of AI. Much of this work focuses on how lawmakers and administrative agencies ought to govern the deployment of these tools (Engstrom, Ho 2020; Engstrom et al. 2020; Solow-Niederman 2020), given that they are already being used by public (Engel, Grgić-Hlačá 2021) and private (Barocas and Selbst 2016) actors alike, producing novel privacy concerns (Ben-Shahar 2023; King et al. 2023) and giving rise to new forms of liability (Henderson, Hashimoto, and Lemley 2023; Lemley and Casey 2019; Volokh 2023). As one highly influential but still maturing species of AI, LLMs stand in need of a concrete empirical evaluation of their legal abilities and their legal risks, of which hallucination is certainly one. We supply that information here.

We also contribute to a growing literature regarding the implications of AI for access to justice. Many members of the legal community rightly regard LLMs as a promising solution to the longstanding barriers to adequate legal representation that millions of *pro se* and under-resourced litigants encounter (Chien et al. 2024; Perlman 2023; Tan, Westermann, and Benyekhlef 2023). Because they are relatively cheap, easy, and quick to use, LLMs might finally be able to deliver on the federal rules’ guarantee of a “just, speedy, and inexpensive” resolution of disputes (Fed. R. Civ. P. 1; Roberts 2023). This potential can only be realized, however, if LLMs actually know the law. Additionally, if the legal knowledge embedded in LLMs is not evenly distributed, the widespread adoption of LLMs might unintentionally worsen rather than eliminate current disparities in the availability of legal services (Draper and Gillibrand 2023; Simshaw 2022). We therefore approach our study of LLMs with an eye toward assessing their ability to truly close the justice gap, examining both their raw hallucination rates as well as any other emergent behaviors that threaten this potential.

Finally, we also contribute to the pressing algorithmic harm literature, which is motivated by the concern that inscrutable algorithms often produce predictions, recommendations, or outputs that are not fairly distributed among individuals or groups (Bar-Gill, Sunstein, and Talgam-Cohen 2023; Gillis and Spiess 2019; Kleinberg et al. 2018; Mayson 2019). In our legal setting, the specific danger is that if LLMs do not properly internalize knowledge about some dimension of the law—if LLMs know California law better than Wyoming law, for example, or decisions by Justice Ketanji Brown Jackson worse than decisions by Justice Antonin Scalia, for another—they will regurgitate a falsely homogeneous sense of the legal landscape to their users, collapsing important legal nuances and perpetuating representational harms. Worse, because LLMs are so-called “foundation” models, their distributional biases, if they exist, may permeate and afflict every downstream version of these models (Bommasani et al. 2022), producing a kind of algorithmic “monoculture” by entrenching one particular notion of the law across a wide range of applications (Creel and Hellman 2022; Kleinberg and Raghavan 2021). Accordingly, it is important for legal scholars to obtain a sense of what the correlates of LLMs’ hallucinations are, in order to address this new and profound opportunity for cascading algorithmic harms.

Our article proceeds as follows. In Section 2, we provide a brief background on LLMs for the non-technical reader and theorize a typology of legal hallucinations. In Section 3, we develop a set of legal knowledge queries that we use to elicit an LLM’s understanding of the law, from simple queries

like whether or not a case exists to more complex queries like asking for a statement of a case's holding or its precedential relationship to another case. In Section 4, we describe our methodological approach, which entails asking these queries for a random sample of cases across each level of the federal judiciary—the US District Courts (USDC), the US Courts of Appeals (USCOA), and the US Supreme Court (SCOTUS)—and evaluating them using four popular LLMs: OpenAI's ChatGPT 4, OpenAI's ChatGPT 3.5, Google's PaLM 2, and Meta's Llama 2.

In Section 5, we present our results. Our findings reveal the widespread occurrence of legal hallucinations: when asked a direct, verifiable question about a randomly selected federal court case, LLMs hallucinate between 58% (ChatGPT 4) and 88% (Llama 2) of the time. However, we also find that LLMs perform better on cases that are newer, more salient, and from more prominent legal jurisdictions, suggesting that the risks of legal monoculture are real. We then investigate two additional potential failure points for LLMs, beyond their raw hallucination rates: (i) their susceptibility to contra-factual bias, i.e., their ability to respond to queries anchored in erroneous legal premises (Sharma et al. 2023; Wei et al. 2023), and (ii) their certainty in their responses, i.e., their self-awareness of their propensity to hallucinate (Azaria and Mitchell 2023; Kadavath et al. 2022; Tian, Mitchell, Zhou, et al. 2023; Xiong et al. 2023; Yin et al. 2023). Our results indicate that not only do LLMs often provide seemingly legitimate but incorrect answers to contrafactual legal questions, they also struggle to accurately gauge their own level of certainty without post-hoc recalibration. Accordingly, in Section 6 we conclude that while LLMs appear to offer a way to make legal information and services more accessible and affordable to all, their present shortcomings—particularly in terms of generating accurate and reliable statements of the law—significantly hinder this objective.

2 BACKGROUND AND THEORY

2.1 What Is a Language Model?

We first provide a brief overview of language models (LMs) for readers who may not necessarily have a deep technical background. LMs can be viewed as functions that map text to text: When a user provides a text input (known as a “prompt”), the model produces a text output (referred to as a “response”). If the prompt takes the form of a question, the response can be understood as an answer to that question. An LM generates its response by selecting the most probable sequence of tokens that follow the prompt's tokens; therefore, it essentially functions as a probability distribution over these tokens.

In this work, we focus on *large* language models (LLMs). The largeness of a language model is a dual reference to its parameter count and the scope of its training corpus: LLMs are models that contain billions of parameters and are trained on vast corpora bordering on the size of the Internet. Because of their incredible size, LLMs can be considered general purpose technologies, with the apparent ability to understand and generate human-like text across a wide range of topics, including medicine, finance, education, retail, and law (Eloundou et al. 2023). In contrast to previous forms of machine learning, however, they seem to excel at these tasks despite not being explicitly trained to perform them (Brown et al. 2020); the “jagged frontier” of their emergent abilities is still being mapped (Dell'Acqua et al. 2023).

We also set forth a more formal definition of an LLM, in order to provide the foundation for the typology of legal hallucinations that we develop in the next subsection. We let an LLM be a function f_τ : prompt \mapsto response, where f_τ operates by sampling responses from a conditional probability distribution that is learned by optimizing over a training corpus hopefully reflective of facts about the world.¹ We use the symbol τ to designate a user-configurable “temperature” parameter that controls the shape of the probability distribution at inference time. When $\tau = 0$, the distribution becomes degenerate and the model's response is theoretically deterministic—the model must always return the most likely response.² Following convention, we refer to this deterministic response as the model's “greedy” response. As τ increases, the distribution becomes more uniform and the model's response becomes more

¹ In reality, language generation in LMs actually happens at the level of tokens, not responses themselves; full model responses are constructed autoregressively by sampling n tokens, one at a time, from a distribution $Pr[x_n | x_1, \dots, x_{n-1}]$. We abstract from these details in this article without loss of generality.

² Non-determinism may persist in practice due to a model's implementation details, e.g., the “mixture of experts” architecture (Chann 2023).

stochastic—the model is free to choose from a variety of candidate responses, all of which become more equally likely to be chosen the higher the temperature is. Thus, increasing the temperature of an LLM is one way to potentially increase its hallucination frequency (Lee 2023). In this article, however, we generally perform our experiments at $\tau = 0$, showing that LLMs hallucinate even under the most conservative sampling conditions.

2.2 The Nature of Legal Hallucinations

LLMs are showing promise on a number of legal research and analysis tasks (Ash et al. 2024; Blair-Stanek, Holzenberger, and Van Durme 2023; Choi et al. 2022; Fei et al. 2023; Guha et al. 2023; Katz et al. 2023; Trozze, Davies, and Kleinberg 2023), but the problem of legal hallucination has so far only been studied in closed-domain applications, such as when a model is used to summarize the content of a given judicial opinion (Deroy, Ghosh, and Ghosh 2023; Feijo and Moreira 2023) or to synthesize provided legal text (Savelka et al. 2023). In this article, by contrast, we examine hallucination in an open-domain setting, i.e., when a model is tasked with providing an accurate answer to an open-ended legal query. This setting approximates the situation of a lawyer or a *pro se* litigant seeking advice from a legal chat interface.

In the context of such question-answering (QA) scenarios, the study of hallucinations in LMs is still in its infancy, even outside the legal field. There is no universally accepted definition or classification of LM hallucinations (Ji et al. 2023; Zhang et al. 2023; van Deemter 2024). However, as Kalai and Vempala (2023) show, LMs that assign a positive probability to every response token must hallucinate at least some of the time. Xu, Jain, and Kankanhalli (2024) agree, arguing that “hallucination is inevitable for any computable LLM, regardless of model architecture, learning algorithms, prompting techniques, or training data.” Therefore, if hallucinations are here to stay, we believe that it is essential for legal scholars to begin to recognize that there are several different ways in which an LLM can generate false information, as not all modes of hallucination are equally concerning for legal professionals. For example, since hallucinations seem likely to give rise to new forms of tort liability (Henderson, Hashimoto, and Lemley 2023), it will be important to differentiate between different types of hallucinations to properly assess the predicate elements of such torts. We supply those theoretical resources here, summarizing our typology of legal hallucinations in Table 1.

First, a model might hallucinate by producing a response that is either unfaithful to or in conflict with the input prompt, a phenomenon canonically referred to as *closed-domain* or *intrinsic* hallucination. This is a major concern in tasks requiring a high degree of accuracy between the response and a long-form input, such as machine translation (Xu et al. 2023) or summarization (Cao et al. 2018). In legal contexts, such inaccuracies would be particularly problematic in activities like summarizing judicial opinions, synthesizing client intake information, drafting legal documents, or extracting key points from an opposing counsel’s brief.

Second, an LLM might also hallucinate by producing a response that either contradicts or does not directly derive from its training corpus. Following Agrawal et al. (2023), we conceptualize this kind of hallucination as one form of *open-domain* or *extrinsic* hallucination. In general, the output of a language model should be logically derivable from the content of its training corpus, regardless of whether the content of the corpus is factually or objectively true.³ In legal settings, this kind of hallucination poses a special challenge to those aiming to fine-tune the kind of general-purpose foundation models that we study in this article with proprietary, in-house work product.⁴ For example, firms might have a

Table 1 Typology of legal hallucinations

Domain	Type of hallucination	Legal example
Closed	Response inconsistency with the prompt	Mischaracterization of an opinion
Open	Response inconsistency with the training corpus	Creative argumentation
	Response inconsistency with the facts of the world	Misstatement of the law

³ For example, if a training corpus consisted of J. K. Rowling’s *Harry Potter* series, we would expect an LLM to produce the sentence “Tom Marvolo Riddle” in response to a query about Voldemort’s real name. However, if the training corpus consisted solely of Jane Austen’s *Pride and Prejudice* (for instance), we would consider this LLM output to be a hallucination—because there would be no basis in the training data for making such a claim about Voldemort.

⁴ For example, this kind of firm-specific fine-tuning is the business model of a prominent legal tech startup, Harvey.ai (Ambrogio 2023).

catalogue of internal research memos, style guides, and so forth, that they want to ensure is reflected in their bespoke LLM's output. At the same time, however, insofar as creativity is valued, certain legal tasks—such as persuasive argumentation—might actually benefit from some lack of strict fidelity to the training corpus; after all, a model that simply parrots exactly the text that it has been trained on could itself be undesirable. As mentioned, creativity can be induced by raising the temperature of the LLM, but responses that are more unpredictable are also those that are more likely to be hallucinations (Lee 2023). Thus, defining the contours of what counts as an *unwanted* hallucination in this specific sense requires value judgements about the balance between fidelity and spontaneity.

Finally, the third way that an LLM can hallucinate is by producing a response that lacks fidelity to the facts of the world, irrespective of how the LLM is trained or prompted (Maynez et al. 2020). We consider this to be another type of open-domain hallucination, with the key concern being “factuality” in relation to the facts of the world (cf. Wittgenstein, 1998 [1921]). In our context, this is perhaps the most alarming type of hallucination, as it can undermine the accuracy required in any legal context where a correct statement of the law is necessary.

2.3 Hallucination Trade-offs

In this article, we investigate only the last kind of hallucination. As mentioned, the first two modes of hallucination are not always problematic in the legal setting: these kinds of hallucinations could actually be somewhat desirable to lawyers if they resulted in generated language that, for example, removed unnecessary information from a given argument (at the expense of being faithful to it) or invented a novel analogy never yet proposed (at the expense of being grounded in the lexicon) (Cao, Dong, and Cheung 2022). However, what a lawyer cannot tolerate is the third kind of hallucination, or factual infidelity between an LLM's response and the controlling legal landscape. In a common law system, where *stare decisis* requires attachment to the “chain” of historical case law (Dworkin 1986), any misstatement of the binding content of that law would make an LLM quickly lose any professional or analytical utility.

Focusing on non-factual hallucinations alone, however, comes with certain trade-offs. One of the advantages of our typology is that it makes clear that it may not always be possible to minimize all modes of hallucination simultaneously; reducing hallucinations of one kind may increase hallucinations of another. For example, if a given prompt contains information that does not conform to facts

Table 2 Hallucination QA task list. Tasks are sorted in order of increasing complexity. Query wording is paraphrased; see the [Online Appendix](#) for exact prompt used. Method column describes the inferential strategy that we use to estimate a hallucination rate for each task: reference-based tasks use known metadata to assess hallucinations, and reference-free tasks use emergent contradictions to assess hallucinations (see Section 4).

Complexity	Task	Query	Method
Low	Existence	Is {case} a real case?	Reference-based
	Court	What court decided {case}?	Reference-based
	Citation	What is the citation for {case}?	Reference-based
	Author	Who wrote the majority opinion in {case}?	Reference-based
Moderate	Disposition	Did {case} affirm or reverse?	Reference-based
	Quotation	What is a quotation from {case}?	Reference-based
	Authority	What is an authority cited in {case}?	Reference-based
	Overruling year	What year was {case} overruled?	Reference-based
High	Doctrinal agreement	Does {case1} agree with {case2}?	Reference-based
	Factual background	What is the factual background of {case}?	Reference-free
	Procedural posture	What is the procedural posture of {case}?	Reference-free
	Subsequent history	What is the subsequent history of {case}?	Reference-free
	Core legal question	What is the core legal question in {case}?	Reference-free
	Central holding	What is the central holding in {case}?	Reference-free

about the world, then ensuring response fidelity with respect to the former would by definition produce infidelity—i.e., hallucination—with respect to the latter. More generally, although fidelity to the prompt is necessary for avoiding *closed-domain* hallucination, there is an important sense in which prioritizing such behavior might actually induce the kind of *open-domain* hallucination that we center in this article.

These trade-offs present unavoidable challenges for prospective users of legal LLMs. When responding to a query, should an LLM be skeptical of its prompt or sycophantic to it? If it has been trained on case law from one jurisdiction, should it enforce adherence to that training corpus even when responding about the law in another jurisdiction? If facts about the world conflict with each other—as legal rules often do—should the LLM preserve that nuance or refrain from introducing information outside the scope of a query? Questions like these are ultimately questions about which kinds of legal hallucinations are more and less preferable, and they are questions whose answers require *both* empirical evidence *and* normative arguments. For example, minimizing fact and training corpus hallucinations (at the expense of prompt hallucinations) might be best for avoiding harm to *pro se* litigants, but the calculus might be reversed for sophisticated lawyers who might be less vulnerable to such behavior. We supply some of the empirics that speak to these dilemmas (see Sections 5.1.6 and 5.2), but stress that the normative considerations are crucial and should be a topic of continued legal hallucination research.

3 PROFILING HALLUCINATIONS USING LEGAL KNOWLEDGE QUERIES

To empirically assess the incidence and correlates of non-factual hallucinations, we adopt a QA framework where the goal is to test an LLM's ability to produce accurate information in response to different kinds of legal queries. We develop fourteen tasks representative of such queries, which we group into three categories in order of increasing complexity and list in Table 2.

3.1 Low Complexity Tasks

In the low complexity category, we ask for information that we consider relatively easy for an LLM to reproduce. The information in this category does not derive from the actual content of a case itself, so it does not require higher-order legal reasoning skills to internalize. Instead, this information is readily available in a case's caption or its syllabus—standard textual locations whose patterns even non-specialized LLMs should be able to recover. We therefore expect LLMs to perform best on these tasks:

3.1.1 Existence

Given the name and citation of a case, state whether the case actually exists or not. This basic evaluation provides preliminary insights into an LLM's knowledge of actual legal cases: if it cannot distinguish real cases from non-existent ones, it probably cannot offer detailed case insights. We use only real cases in our prompts, so affirming their existence is the correct answer.⁵

3.1.2 Court

Given the name and citation of a case, supply the name of the court that ruled on it. This task assesses an LLM's knowledge about legal jurisdictions, an important building block of a case's precedential value. We perform this task across the three different levels of the federal judiciary. Importantly, we note that each level of the judiciary has a different reporter, or the series of volumes that opinions are published in. This is relevant because the reporter is included in the citation that we provide to the LLM, essentially revealing the level of the hierarchy that an opinion is from. All and only SCOTUS cases are published in the *US Reports*. Opinions from the USCOA are published in the *Federal Reporter*, and USDC cases are published in the *Federal Supplement*. Because of this, we expect this task to be more difficult as we descend the hierarchy of courts. There is only one court associated with the US reporter, but 13 associated with the *Federal Reporter*, and 94 associated with the *Federal Supplement*. For USCOA cases, we require the name of the specific circuit court, and for USDC cases, we require the name of the specific district court.

3.1.3 Citation

Given a case name, supply the Bluebook citation of the case. This query tests an LLM's ability to associate a given dispute with its official record in a reporter volume at a particular page, which is the key way in

⁵ In the [Online Appendix](#), we experiment with using fake cases as well.

which different opinions reference and link to each other. For USCOA cases, we further specify that we want the citation for the circuit court opinion, and for USDC cases, we further specify that we want the citation for the district court opinion. We test for citation equality using *eyecite* (Cushman, Dahl, and Lissner 2021).

3.1.4 Author

Given the name and citation of a case, supply the name of the opinion author. This query tests an LLM's ability to associate a given case with a particular judge, which is important for contextualizing a case in the broader jurisprudential landscape. For SCOTUS and USCOA cases, we further specify that we want the name of the majority opinion author. We accept a fuzzy match of the opinion author's name as accurate.

3.2 Moderate Complexity Tasks

Next, in the moderate complexity category, we start to require an LLM to evince knowledge of actual legal opinions themselves. To answer the queries in this category, an LLM must know something about a case's substantive content; these queries seek information that must be collated from idiosyncratic portions of its text. Of course, a database-augmented LLM might still be able to retrieve some of this information without ever actually internalizing the content of a case, but we expect this text-based knowledge to be less available than the information described in the low complexity category. Specifically, we ask for the following information:

3.2.1 Disposition

Given a case name and its citation, state whether the court affirmed or reversed the lower court. This query tests an LLM's knowledge of how the court resolved the instant appeal confronting the parties in the case, which is the first step for determining the holding that is created by the case. Though this is essentially a binary classification task where we accept correct "affirm" or "reverse" labels as accurate, we consider this task to still be probative of hallucinations because producing the wrong label is still a misstatement of the law. We filter out all ambiguous dispositions (e.g., reversals in part) and we do not ask this query of USDC cases because district courts are courts of original jurisdiction.⁶

3.2.2 Quotation

Given a case name and its citation, supply any quotation from the opinion. This query tests an LLM's ability to produce some portion of an opinion's text verbatim, which is an important feature for lawyers seeking to use a case to stand for a specific proposition. Normally, such memorization is considered an undesirable property of LLMs (Carlini et al. 2022), but in this legal application it is actually desirable behavior. We accept any fuzzy string of characters appearing in the majority opinion as accurate.

3.2.3 Authority

Given a case name and its citation, supply a case that is cited in the opinion. This query probes an LLM's understanding of the chain of precedential authority that supports a given opinion. We do not distinguish between positive and negative citations for this task; we accept any precedent cited in any way in the text of the majority opinion as accurate. We extract and match citations on their volumes, reporters, and pages using *eyecite* (Cushman, Dahl, and Lissner 2021).

3.2.4 Overruling Year

Given a case name and its citation, supply the year that it was overruled. This query tests an LLM's ability to recognize when a given case has been subsequently altered, which is crucial information for lawyer seeking to determine whether a given precedent is still good law or not. This task is the most complicated in this category because it requires the LLM to draw connections between multiple areas of the case space. We accept only the exact year of overruling as accurate, and we limit this task to only those SCOTUS cases that have been explicitly overruled ($n = 279$).⁷

⁶ While it is possible for some administrative agency decisions to be appealed to a district court, this occurs infrequently enough that we choose not to ask for case disposition at the district court level.

⁷ In Section 5.2, we experiment with cases that have never been overruled as well.

3.3 High Complexity Tasks

Finally, in the high complexity category, we seek answers to tasks that *both* presuppose legal reasoning skills (unlike the low complexity tasks) *and* are not readily available in existing legal databases like WestLaw or Lexis (unlike the moderate complexity tasks). These tasks all require an LLM to synthesize core legal information out of unstructured legal prose—information that is frequently the topic of deeper legal research. In Section 4.3, we explain how we test LLMs' knowledge of some of these more complex facts without necessarily having access to the ground-truth answers ourselves:

3.3.1 Doctrinal Agreement

Given two case names and their citations, state whether they agree or disagree with each other. This query requires an LLM to show knowledge of the precedential relationship between two different cases, information that is essential for higher-order legal reasoning. We use Shepard's treatment codes as a basis for constructing this task, filtering out all ambiguous citation treatments (e.g., neutral treatments) and coarsening the unambiguous codes into "agree" and "disagree" labels that we accept as accurate. For this task, we use a relatively balanced dataset of 2,839 citing-cited case pairs coded as "agree," and 2,161 citing-cited case pairs coded as "disagree." This task is limited to SCOTUS cases, as our underlying dataset only contains thorough Shepard's data for citations to the Supreme Court.

3.3.2 Factual Background

Given a case name and its citation, supply its factual background. This query tests an LLM's understanding of the concrete fact pattern underlying a case, which is helpful in assessing the relevance of the case to current research and in drawing parallels with other cases.

3.3.3 Procedural Posture

Given a case name and its citation, supply its procedural posture. This query tests an LLM's understanding of how and why a case has arrived at a particular court, which aids in understanding the precise question presented and standard of review applicable.

3.3.4 Subsequent History

Given a case name and its citation, supply its subsequent procedural history, if any. This query tests an LLM's knowledge of any other related proceedings that concern the given case after a particular decision, which is information that can change or clarify the legal significance of the case.

3.3.5 Core Legal Question

Given a case name and its citation, supply the core legal question at issue. This query tests an LLM's ability to pinpoint the main issue or issues that a court is addressing in a case, which is the most important factor in assessing whether a case is apposite or not.

3.3.6 Central Holding

Given a case name and its citation, supply its central holding. This query tests an LLM's knowledge of the legal principle that a given case stands for, i.e., the precedent that future cases will rely upon or distinguish from. Articulating the holding of a case is crucial for legal analysis and argumentation and is the most complex task that we evaluate.

4 EXPERIMENTAL DESIGN

4.1 Data Construction

We aim to profile hallucination rates across several legally salient dimensions, including hierarchy, jurisdiction, time, and case prominence. Thus, we construct our test data with an eye toward making statistical inferences on these covariates.

We begin with the universe of case law from each level of the federal judicial hierarchy—namely, SCOTUS, USCOA, and USDC—that has been published in the volumes of the *U.S. Reports*, the *Federal Reporter*, and the *Federal Supplement*. To ensure balance over time and place, we then perform stratified random sampling using year strata for the SCOTUS cases, circuit-year strata for the USCOA

cases, and state-year strata for the USDC cases. We draw 5,000 cases from each level of the judiciary. Finally, to generate ground-truth answers for our reference-based queries (Section 4.2), we merge these units with metadata obtained from the [Caselaw Access Project \(2023\)](#), the Supreme Court Database ([Spaeth et al. 2022](#)), the Appeals Courts Database Project ([Songer 2008](#); [Kuersten and Haire 2011](#)), the Library of Congress ([Congress.gov 2023](#)), and Shepard's Citations ([Fowler et al. 2007](#); [Black and Spriggs 2013](#)).⁸

4.2 Reference-Based Querying

The most straightforward way to study hallucinations in the open-domain setting is to use a test oracle—or an external *reference*—to detect and adjudge non-factual responses ([Lin, Hilton, and Evans 2022](#); [Lee et al. 2023](#); [J. Li et al. 2023](#)). Such oracles are usually difficult and costly to construct ([Krishna, Roy, and Iyyer 2021](#)), but we use the tabular metadata described in Section 4.1 to develop ours. Our design exploits the fact that while LLMs are known to have been trained on the raw text of American case law, which is in the public domain ([Henderson et al. 2022](#)), they have likely not been trained on these cases' attendant metadata, which exist separately from the cases' textual content and which we have aggregated from disparate sources.

These metadata enable us to construct reference-based queries for the first nine of our tasks (Table 2). These queries take the form of N question-and-answer triples (prompt, response, response'), where prompt is a case-specific question, response is the LLM's greedy answer retrieved from calling , and response' is the known ground-truth answer.⁹ Our estimand of interest for each task is the population-level hallucination rate π , which we estimate by averaging over the N sampled queries:

$$\pi = \hat{\pi} = \frac{1}{N} \sum 1[\text{response} \neq \text{response}'] \quad (1)$$

Occasionally, an LLM will produce a response that is neither a hallucination nor a correct answer, but rather an explicit *abstention* from answering the question. For example, the LLM might admit that it does not know the answer or demur that it is unable to provide the answer for some reason, perhaps due to safety concerns. In these instances, we nevertheless count the response as a non-hallucination, on the theory that an LLM cannot hallucinate when it is affirmatively abstaining from responding ([Feng et al. 2024](#)). We document the frequency of these abstentions in the [Online Appendix](#), but they are generally rare and do not substantively affect our findings.

4.3 Reference-Free Querying

Reference-based querying lets us directly recover our population parameter of interest, but two problems limit the effectiveness of the approach. First, we are restricted to asking questions for which digestible metadata exist and a clear answer has been recorded, which rules out many more complex inquiries. Second, precisely because these queries can be answered with tabular data, legal database-augmented LLMs ([Cui et al. 2023](#); [Savelka et al. 2023](#)) are likely to soon solve or at least mask hallucinated responses to these queries ([Peng et al. 2023](#); [Shuster et al. 2021](#)).

To test the tasks that cannot be easily verified against an external legal database, we employ reference-free querying instead, which detects hallucinations by exploiting the stochastic behavior of LLMs at higher temperatures ([Agrawal et al. 2023](#); [Manakul, Liusie, and Gales 2023](#); [Min et al. 2023](#)). This approach is rooted in the theory that hallucinations are more likely to originate in flat probability distributions with higher next-token uncertainties, whereas factual answers should always have a high probability of being the generated response given a prompt. Thus, by repeatedly querying an LLM at a non-greedy temperature, we can estimate the model's hallucination rate by examining its self-consistency—factual responses should not change, but hallucinated ones will.

Most reference-free approaches implicitly assume that the LLM is calibrated, i.e., that there is indeed some correlation between its self-consistency and its propensity to hallucinate. For reasons that we discuss in Section 5.3, we are unwilling to make this assumption in our legal setting. We therefore adopt a slightly different implementation that is still reference-free, but only requires *contradiction*, not consistency ([Mündler et al. 2023](#)). Specifically, for our final five tasks (Table 2), we construct reference-free

⁸ More information about how we use these metadata to construct each query is available in the [Online Appendix](#).

⁹ Recall from Section 2.1 that the j_0 notation represents performing inference with the LLM at temperature zero—i.e., under its deterministic behavior.

queries in the form of N question-and-answer triples (prompt, response⁽¹⁾, response⁽²⁾), where prompt is the question, response⁽¹⁾ is one LLM answer retrieved by calling $f_1(\cdot)$ once, and response⁽²⁾ is another LLM answer retrieved by calling $f_1(\cdot)$ again.¹⁰ Detecting a hallucination then amounts to detecting a logical contradiction between the two stochastic answers: any such contradiction guarantees non-factuality, because two contradictory answers cannot both be correct.

To identify these contradictions at scale, we feed both answers into GPT 4 and ask it for its assessment. This technique does not assume anything about $f_1(\cdot)$'s calibration—it just requires that GPT 4 possess logical reasoning skills sufficient to compare $f_1(\cdot)$'s two responses and accurately label them as contradictory as not. To justify this reliance on GPT 4, we manually label a portion of the reference-free responses ourselves and conduct an intercoder reliability analysis to ensure that GPT 4 is indeed able to perform this task. Full information about our procedure and a validity check is provided in the [Online Appendix](#). (We find that GPT 4's reliability is comparable to human labeling of contradictions.)

An important caveat of this approach is that it only allows us to establish a *lower bound* on the hallucination rate for our reference-free queries:

$$\pi \geq \hat{\pi} = \frac{1}{N} \sum 1[\text{response}^{(1)} \neq \text{response}^{(2)}] \quad (2)$$

Although self-contradiction guarantees hallucination, the inverse does not hold: two answers may be logically consonant but still lack fidelity to the law. Because we are unwilling to assume calibration, we accept this inferential limitation, but, as we show below, even the lower bounds on hallucination rates are quite high and informative.

4.4 Models

We perform our experiments using four popular, state-of-the-art, off-the-shelf LLMs:

1. OpenAI's ChatGPT 4 (gpt-4-1106-preview, [OpenAI 2023a](#)),
2. OpenAI's ChatGPT 3.5 (gpt-3.5-turbo-0613, [2023b](#)),
3. Google's PaLM 2 (text-bison-001, [Anil et al. 2023](#)), and
4. Meta's Llama 2 (Llama-2-13b-chat-hf, [Touvron et al. 2023](#)).

We run each query under both “zero-shot” and “three-shot” prompting setups. In the zero-shot setup, we simply ask the LLM about the given case directly, whereas in the three-shot setup, we prepend several example questions and responses to give the LLM an opportunity to perform in-context learning ([Brown et al. 2020](#)). We provide the full text of the prompts we use for each query, along with the few-shot examples, in the [Online Appendix](#). In total, we execute more than 800,000 queries—200,000+ per LLM—and we share our raw API calls and model responses in the replication materials accompanying this article.

5 Results

We begin by presenting our main results profiling LLMs' hallucination rates, which cut to the core of popular concerns over LLMs' suitability for legal applications (Section 5.1). Then, after showing that hallucinations are generally widespread, and highlighting the correlates of LLMs' hallucination rates, we turn to two additional challenges that threaten LLMs' utility for legal adoption: (i) their susceptibility to contra-factual bias, i.e., their ability to handle queries based on mistaken legal premises (Section 5.2), and (ii) their certainty in their responses, i.e., their self-awareness of their propensity to hallucinate (Section 5.3).

5.1 Hallucination Rates and Heterogeneity

Tables 3–5 report our estimated hallucination rates and their standard errors for each category of our tasks. We find that hallucinations vary with the substantive complexity of the task (Section 5.1.1), the hierarchical level of the court (Section 5.1.2), the jurisdictional location of the court (Section 5.1.3), the

¹⁰ Recall from Section 2.1 that the $f_1(\cdot)$ notation represents performing inference with the LLM at temperature one—i.e., with some degree of stochasticity.

prominence of the case (Section 5.1.4), the year the case was decided (Section 5.1.5), and the LLM queried (Section 5.1.6). We do not find substantial differences between zero-shot and few-shot prompting, so we focus our discussion on the few-shot results alone.

5.1.1 Hallucinations Vary by Task Complexity

As we hypothesized in Section 3, we first observe that hallucinations increase with the complexity of the legal research task at issue, which we visualize in Figure 1. Starting with the low complexity category (Table 3), the LLMs perform best on the simple **Existence** task, though this is in part driven by their tendency to always answer “yes” when asked about the existence of any case. (In the Online Appendix we demonstrate this problem by asking about the existence of fake cases instead.) The models begin to




Figure 1. Relationship between task complexity and mean hallucination rate. Higher values indicate a greater likelihood of factually incorrect LLM responses. High complexity tasks include several reference-free tasks, so those reported hallucination rates are lower bounds on the true rates. Contra-factual tasks and the doctrinal agreement high complexity task are excluded from this comparison.




Figure 2. Relationship between judicial hierarchy and mean hallucination rate, all reference-based tasks pooled. Hallucination rates are higher for lower levels of the federal judiciary.

Table 3 Hallucination rates across levels of the federal judiciary (low complexity tasks)

Task	Prompt	SCOTUS (1794-2015; n = 5,000)				USCOA (1895-2019; n = 5,000)				USDC (1932-2019; n = 5,000)			
		GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2
Existence	Zero-shot	0.204 (0.006)	0.04 (0.001)	0.054 (0.003)	0.303 (0.006)	0.237 (0.006)	0.003 (0.001)	0.025 (0.002)	0.157 (0.005)	0.174 (0.005)	0.001 (0.001)	0.016 (0.002)	0.24 (0.006)
	Few-shot	0.181 (0.005)	0.029 (0.002)	0.029 (0.002)	1.000 (0.000)	0.129 (0.005)	0.018 (0.002)	0.005 (0.001)	1.000 (0.000)	0.048 (0.003)	0.004 (0.001)	0.006 (0.001)	1.000 (0.000)
Court	Zero-shot	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.003 (0.001)	0.490 (0.007)	0.645 (0.007)	0.703 (0.006)	0.700 (0.006)	0.829 (0.005)	0.815 (0.005)	0.839 (0.005)	0.815 (0.005)
	Few-shot	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.005 (0.001)	0.491 (0.007)	0.641 (0.007)	0.703 (0.006)	0.679 (0.007)	0.831 (0.005)	0.87 (0.005)	0.842 (0.005)	0.87 (0.005)
Citation	Zero-shot	0.621 (0.007)	0.684 (0.007)	0.906 (0.004)	0.941 (0.003)	0.727 (0.006)	0.754 (0.006)	1.000 (0.000)	0.999 (0.000)	0.61 (0.007)	0.702 (0.006)	1.000 (0.000)	1.000 (0.000)
	Few-shot	0.689 (0.007)	0.641 (0.007)	0.833 (0.005)	0.953 (0.003)	0.468 (0.007)	0.92 (0.004)	0.999 (0.001)	1.000 (0.000)	0.270 (0.006)	0.942 (0.003)	0.999 (0.000)	1.000 (0.000)
Author	Zero-shot	0.799 (0.006)	0.796 (0.006)	0.816 (0.005)	0.884 (0.005)	0.954 (0.003)	0.965 (0.003)	0.988 (0.002)	0.991 (0.001)	0.922 (0.004)	0.911 (0.004)	0.988 (0.002)	0.987 (0.002)
	Few-shot	0.799 (0.006)	0.83 (0.005)	0.859 (0.005)	0.881 (0.005)	0.962 (0.003)	0.967 (0.003)	0.988 (0.002)	0.993 (0.001)	0.921 (0.004)	0.941 (0.003)	0.984 (0.002)	0.987 (0.002)

Note: Table reports estimated hallucination rates. Standard errors are shown in parentheses.

struggle more when prompted for information about a case's **Court**, **Citation**, or **Author**. Hallucinations then surge among the moderate complexity tasks (Table 4), all of which require the LLMs to evince knowledge of the actual content of a legal opinion. We note that these results are not just a product of different evaluation metrics: although the **Quotation** task, for example, requires near-word reproduction of particular sentences and phrases to be judged correctly, the **Disposition** task simply asks for binary responses from the model. Yet, the LLMs hallucinate widely in both setups.

The results for the high complexity tasks (Table 5) confirm this general pattern of poor performance. Starting with **Doctrinal agreement**, recall that this query asks the LLM to make an analogical judgment about the precedential relationship between two given cases, for which we have ground-truth labels from Shepard's treatment codes. Because this is another binary classification task, the LLMs' hallucination rates on this task—near 0.5—represent little improvement over random guessing, and are actually sometimes worse. This suggests that LLMs know little about substantive legal doctrine, calling into question their ability to accurately assist lawyers in more realistic, applied settings.

The remaining tasks in the high complexity category amplify these concerns, but it is important to keep in mind that the hallucination rates that we report for these tasks are only *lower bounds* on the true rates, as these tasks are evaluated using our reference-free method (Section 4.3). To provide some context for these bounds, we note that in a similar self-contradiction setup, Mündler et al. (2023) found that GPT 3.5 hallucinated about 14.3% of the time on general QA queries. On our legal QA queries, GPT 3.5 and our other LLMs far surpass this baseline rate—and it is possible that the true hallucination rate is even higher.

For example, we find that even on the easier reference-free tasks—**Factual background** and **Procedural posture**—our LLMs hallucinate at least 49% of the time. Performance degrades further on the most complex **Core legal question** and **Central holding** tasks, with hallucinations arising in response to at least 59% and 63% of our queries, respectively. Hallucinations are lowest among GPT 4 responses to the **Subsequent history** task at the SCOTUS level, but this is because the model simply tends to state that the litigation concluded with the Supreme Court decision. This may not actually be correct—many Supreme Court cases result in a remand and have additional procedural history in lower courts. However, we are unable to capture this kind of mistake, as our methodology only permits us to identify hallucinations where the model contradicts itself. We are not able to capture repeated incorrect answers as instances of hallucination, meaning that our estimate of hallucination in the SCOTUS **Subsequent history** task is likely to understate the rate of hallucination by a larger margin than other tasks.

Taken together, these results invite skepticism about LLMs' true knowledge of the law. Our reference-free tasks, in particular, raise serious doubts about LLMs' knowledge of substantive aspects of American case law—the very knowledge that attorneys must often synthesize themselves, instead of merely looking up in a database.

5.1.2 Hallucinations Vary by Court

We next examine trends by hierarchy, exploring LLMs' abilities to restate the case law of the three different levels of the federal judiciary. We find that across all tasks and all LLMs, hallucinations are lowest in the highest levels of the judiciary, and vice-versa (Figure 2). Thus, our LLMs perform best on tasks at the SCOTUS level, worse on tasks at the USCOA level, and worst on tasks at the USDC level. These results are encouraging insofar as it is important for LLMs to be knowledgeable about the most authoritative and wide-ranging precedents, but discouraging insofar as they suggest that LLMs are not well attuned to localized legal knowledge. After all, the vast majority of litigants do not appear before the Supreme Court and may benefit more from knowledge that is tailored to their home district court—their court of first appearance.

5.1.3 Hallucinations Vary by Jurisdiction

To better understand the relationship between different courts and hallucinations, we next zoom in on the middle level of the judicial hierarchy—the Courts of Appeals—and examine horizontal heterogeneity across the circuits.¹¹ Figure 3 depicts these results geographically, showing lower hallucination

¹¹ Because not all Courts of Appeals were created at the same time, for parity in comparison here we exclude from our results cases decided before 1982, the year the youngest circuit—the Federal Circuit—was created. We report the full, non-truncated results in the [Online Appendix](#), which are largely consistent with these post-1981 results.

Table 4 Hallucination rates across levels of the federal judiciary (moderate complexity tasks)

Task	Prompt	SCOTUS (1794–2015; n = 5,000)					USCOA (1895–2019; n = 5,000)					USDC (1932–2019; n = 5,000)						
		GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2	
Disposition	Zero-shot	0.399 (0.007)	0.499 (0.001)	0.500 (0.000)	0.536 (0.007)	0.494 (0.005)	0.500 (0.000)	0.500 (0.000)	0.493 (0.008)									
	Few-shot	0.452 (0.007)	0.496 (0.002)	0.501 (0.002)	0.502 (0.004)	0.498 (0.006)	0.489 (0.007)	0.501 (0.001)	0.501 (0.002)									
Quotation	Zero-shot	0.312 (0.007)	0.229 (0.006)	0.993 (0.001)	0.92 (0.004)	0.001 (0.000)	0.000 (0.000)	0.999 (0.000)	0.993 (0.001)									
	Few-shot	0.854 (0.000)	1.000 (0.000)	0.993 (0.001)	0.992 (0.001)	0.637 (0.007)	1.000 (0.000)	0.997 (0.001)	1.000 (0.000)									
Authority	Zero-shot	0.922 (0.004)	0.937 (0.003)	0.985 (0.002)	0.991 (0.001)	0.934 (0.004)	0.978 (0.002)	0.997 (0.001)	0.999 (0.001)									
	Few-shot	0.828 (0.005)	0.916 (0.004)	0.953 (0.003)	0.993 (0.001)	0.958 (0.003)	0.976 (0.002)	0.993 (0.001)	0.999 (0.000)									
Overruling years	Zero-shot	0.81 (0.025)	0.919 (0.017)	0.858 (0.022)	0.972 (0.011)													
	Few-shot	0.725 (0.028)	0.976 (0.010)	0.87 (0.021)	0.984 (0.008)													

^a 1810–2022 (n = 279).

Note: Table reports estimated hallucination rates. Standard errors are shown in parentheses.

Table 5 Hallucination rates across levels of the federal judiciary (high complexity tasks)

Task	Prompt	SCOTUS (1794–2015; n = 100)				USCOA (1895–2019; n = 100)				USDC (1932–2019; n = 100)			
		GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2
Doctrinal agreement	Zero-shot	0.461 (0.005)	0.500 (0.000)	0.466 (0.005)	0.500 (0.000)								
	Few-shot	0.449 (0.007)	0.458 (0.004)	0.453 (0.006)	0.500 (0.000)								
Factual background	Zero-shot	0.46 (0.05)	0.70 (0.046)	0.95 (0.022)	0.83 (0.038)	0.75 (0.043)	0.81 (0.039)	0.88 (0.032)	0.88 (0.032)	0.71 (0.045)	0.67 (0.047)	0.88 (0.032)	0.88 (0.032)
	Zero-shot	0.47 (0.05)	0.65 (0.048)	0.74 (0.044)	0.85 (0.036)	0.66 (0.047)	0.6 (0.049)	0.68 (0.047)	0.83 (0.038)	0.73 (0.044)	0.55 (0.05)	0.81 (0.039)	0.89 (0.031)
Procedural posture	Zero-shot	0.08 (0.027)	0.22 (0.041)	0.81 (0.039)	0.77 (0.042)	0.45 (0.050)	0.36 (0.048)	0.75 (0.043)	0.62 (0.049)	0.48 (0.05)	0.3 (0.046)	0.73 (0.044)	0.67 (0.047)
	Zero-shot	0.57 (0.05)	0.86 (0.035)	0.87 (0.034)	0.92 (0.027)	0.81 (0.039)	0.88 (0.032)	0.95 (0.022)	0.96 (0.02)	0.76 (0.043)	0.72 (0.045)	0.92 (0.027)	0.89 (0.031)
Core legal question	Zero-shot	0.6 (0.049)	0.73 (0.044)	0.78 (0.041)	0.92 (0.027)	0.83 (0.038)	0.87 (0.034)	0.95 (0.022)	0.95 (0.022)	0.77 (0.042)	0.73 (0.044)	0.93 (0.026)	0.84 (0.037)
	Zero-shot												

^a 1796–2005 (n = 5,000).

Note: Table reports estimated hallucination rates. For all tasks except doctrinal agreement, this rate is only a lower bound on the true population rate. Standard errors are shown in parentheses.

rates in lighter colors and higher rates in darker colors. Pooling our tasks and models together, we see the best performance in the Ninth Circuit (comprising California and adjacent states in yellow), the Second Circuit (comprising New York and adjacent states in soft green), the Third Circuit (comprising Pennsylvania and adjacent states in soft green), and the First Circuit (comprising Maine and adjacent states in soft green). By contrast, performance tends to be worst in the circuits in the geographic center of the country.

These results confirm popular intuitions about the influential role that the Second, Third, and Ninth Circuits play in the American legal system. Because it encompasses New York City, the Second Circuit has traditionally had a significant impact on financial and corporate law, and many landmark decisions in securities law, antitrust, and business litigation have come from this court. The Third Circuit enjoys similar influence in the corporate law domain owing to Delaware's status as the legal home for many corporations. Finally, the Ninth Circuit handles more cases than any other federal appellate court, and often issues rulings that advance progressive positions that lead to disproportionate review by the Supreme Court.

Perhaps surprisingly, however, our results stand in tension with received wisdom about the D.C. Circuit, which is generally thought to be the *most* influential appellate division. In our tasks, our LLMs actually perform worst on this circuit. This counterintuitive finding is one example of the way that unanticipated biases might trouble the reliance on LLMs in practice.

5.1.4 Hallucinations Vary by Case Prominence

To probe the role of legal prominence more directly, we move to SCOTUS-level results next, examining the relationship between case importance and hallucinations. To measure case prominence within this single level of the judiciary, we use the Caselaw Access Project's PageRank percentile scores, a metric of citation network centrality that captures the general legal and political prominence of a case.

We find that case prominence is negatively correlated with hallucination, reaffirming our results from above (Figure 4). However, we also note that a sharp slope change occurs around the 90th prominence percentile in the GPT 4, GPT 3.5, and PaLM 2 models. This suggests that the bias of these LLMs—but not Llama 2—may be skewed even more toward the most well-known decisions of the American legal system, even within the SCOTUS level.

5.1.5 Hallucinations Vary by Case Year

Because case law develops in virtue of new decisions building on old ones over time, the age of a case may be another useful predictor of hallucination. Examining this relationship at the SCOTUS level in Figure 5, we find a non-linear correlation between hallucination and age: hallucinations are most common among the Supreme Court's oldest and newest cases, and least common among its post-war Warren Court cases (1953–1969). This result suggests another important limitation on LLMs' legal knowledge that users should be aware of: LLMs' peak performance may lag several years behind the current state of the doctrine, and LLMs may fail to internalize case law that is very old but still applicable and relevant law.

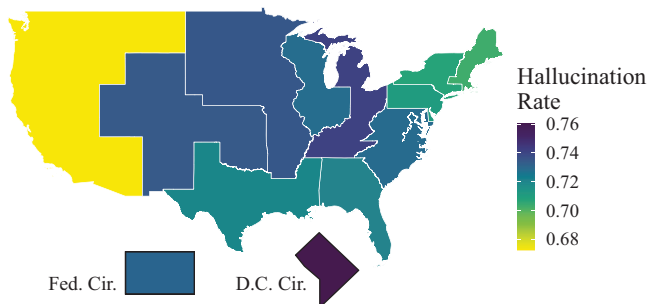


Figure 3. Relationship between USCOA jurisdiction and mean hallucination rate, all reference-based US- COA tasks and models pooled, post-1981 cases only. LLM performance is strongest in jurisdictions that are commonly perceived to play a more influential role.

5.1.6 Hallucinations Vary by LLM

Finally, we also partition our results by the LLM itself and compare across models. We find that not all LLMs are equal: as expected, GPT 4 performs best overall, followed by GPT 3.5, followed by PaLM 2, followed by Llama 2 (Figure 6).

We also discover tendencies towards different inductive biases, or the predisposition of an LLM to generate certain outputs more frequently than others. In Figure 7, we highlight one of these biases for our SCOTUS-level **Author** task, which asks the LLM to supply the name of the justice who authored the majority opinion in the given case. Each LLM we test has slightly different inductive preferences;




Figure 4. Relationship between SCOTUS case prominence (measured by PageRank percentile) and mean hallucination rate, all SCOTUS tasks pooled. Hallucinations decline sharply as case prominence passes the 90th percentile, meaning that LLMs are more likely to respond with accurate information about prominent cases.




Figure 5. Relationship between SCOTUS case decision year and mean hallucination rate, all SCOTUS tasks pooled. LLMs are most likely to respond with accurate information in cases from the latter half of the 20th century, struggling on very old and very new cases.

some err towards the most recognizable justices, but others are a little more difficult to explain. For example, Llama 2 disproportionately favors Justice Story—an influential jurist who authored the famous *Amistad* opinion, among others—whereas PaLM 2 prefers Justice McLean—also an important jurist, but one more known for his dissents than his majority opinions, such as his dissent in the infamous *Dred Scott* case. Across the board, all our LLMs tend to overstate the true prevalence of justices at a higher magnitude than they understate them, as indicated by the greater dispersion of the points above the $y = x$ line in Figure 7.

These biases demonstrate one way that LLMs inevitably encounter the kind of hallucination trade-off that we discuss in Section 2.3. If the inductive bias that an LLM learns from its training corpus is not




Figure 6. Hallucination rates by LLM, all reference-based tasks pooled. Hallucinations are common across all LLMs when they are asked a direct, verifiable question about a federal court case, but GPT 4 performs best overall.

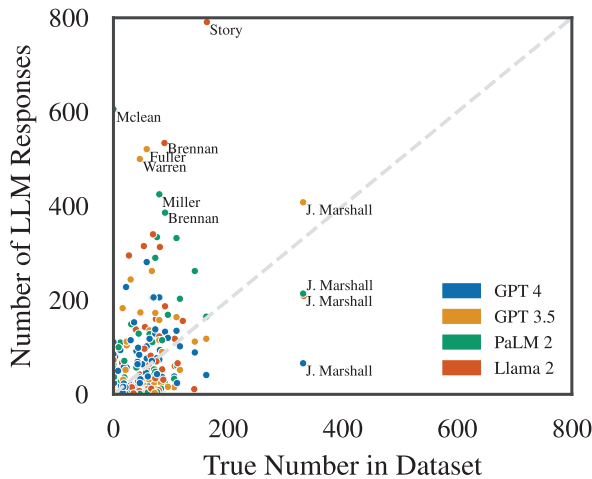


Figure 7. Number of times each justice is stated to be the author of a SCOTUS case versus the actual number of cases authored by each justice in our time period-stratified dataset. A small number of justices are disproportionately represented in LLM responses.

well-aligned with the true distribution of facts about the world, then the LLM is likely to make systematic errors when queried about those facts. Moreover, the persistence of inductive biases also increases the risk of LLMs instantiating a kind of legal monoculture (Kleinberg, Raghavan 2021). Instead of accurately restating the full variation of the law, LLMs may simply regurgitate information from a few prominent members of the response set that they have been trained on, flattening legal nuance and producing a falsely homogenous sense of the legal landscape.

5.2 Contra-factual Bias

We now turn to the first of two potential failure points that we seek to examine for LLMs performing legal tasks, beyond their sheer propensity to hallucinate: their bias toward accepting legal premises that are not anchored in reality and answering queries accordingly. We view this behavior as a particular kind of model sycophancy (the tendency of an LLM to agree with a user's preferences or beliefs, even when the LLM would reject the belief as wrong without the user's prompting; Sharma et al. 2023; Wei et al. 2023) or general cognitive error (Tversky and Kahneman 1974; Jones, Steinhardt 2022; Suri et al. 2023).

This bias poses a subtle but pernicious challenge to those aiming to use LLMs for legal research. When a researcher is learning about a topic, they are not only unsure about the answer, they are also often unsure about the question they are asking as well. Worse, they might not even be aware of any defects in their query; research by its nature ventures into the realm of "unknown unknowns" (Luft and Ingham 1955). This is especially true for unsophisticated *pro se* litigants, or those without much legal training to begin with. Relying on an LLM for legal research, they might inadvertently submit a question premised on non-factual legal information or folk wisdom about the law. As discussed in Section 2.3, this then forces a trade-off: if the LLM is too intent on minimizing prompt hallucinations, it runs the risk of simply accepting the user's misconception as true and producing a factual hallucination instead.

To test whether this risk is real in the legal setting, we evaluate two modified versions of our reference-based queries, but with premises that are false by construction. Specifically, we ask the LLMs to (i) provide information about an author's dissenting opinion in an appellate case in which they did not in fact dissent and (ii) furnish the year that a SCOTUS case that has never been overruled was overruled. In both cases, we consider failing to provide the requested information an acceptable answer; any uncritical answering of the prompt is treated as a hallucination.

Table 6 reports the results of this experiment and Figure 8 summarizes them by LLM. In general, LLMs seem to suffer from contra-factual bias on these legal information tasks. As in the raw hallucination tasks, contra-factual bias hallucinations are higher in lower levels of the judiciary. Substantively, they are also greatest for the question with a false overruling premise, possibly reflecting the increased complexity of the question asked.

Llama 2 performs exceptionally well, demonstrating little contra-factual hallucination. However, this success is linked to a different kind of hallucination—in many false dissent examples, for instance, Llama 2 often states that the case or justice does not exist at all. (In reality, all of our false dissent examples were created with real cases and real justices—just justices who did not author a dissent for the case.) Under our metrics for contra-factual hallucination, we choose to record these examples as successful rejections of the premise. The kind of error that Llama 2 makes here is already measured in its poor performance on other tasks, especially **Existence**.

5.3 Model Calibration

The second potential hazard that we investigate is model calibration, or the ability of LLMs to "know what they know." Ideally, a well-calibrated model would be confident in its factual responses, and not confident in its hallucinated ones (Azaria and Mitchell 2023; Kadavath et al. 2022; Tian, Mitchell, Zhou, et al. 2023; Xiong et al. 2023; Yin et al. 2023). If this property held for legal queries, users would be able to adjust their expectations accordingly and could theoretically learn to trust the LLM when it is confident, and learn to be more skeptical when it is not (Zhang, Liao, and Bellamy 2020). Even more importantly, if an LLM knew when it was likely to be hallucinating, the hallucination problem could be in principle solvable through some form of reinforcement learning from human feedback (RLHF) or fine-tuning, with unconfident answers simply being suppressed (Tian, Mitchell, Yao, et al. 2023).

To study our LLMs' calibration on legal queries, we estimate the expected calibration error (ECE) for each of our tasks. We describe our estimation strategy in full in the [Online Appendix](#), but, intuitively,

Table 6 Hallucination rates across levels of the federal judiciary (contra-factual tasks)

Task	Prompt	SCOTUS (1794–2015; n = 1,000)				USCOA (1895–2019; n = 1,000)			
		GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2
False dissent premise	Zero-shot	0.691 (0.015)	0.338 (0.015)	0.990 (0.003)	0.000 (0.000)	0.842 (0.012)	0.408 (0.016)	0.983 (0.004)	0.021 (0.005)
False overruling premise	Zero-shot	0.531 (0.016)	0.821 (0.012)	1.000 (0.000)	0.027 (0.005)				

Note: Table reports estimated hallucination rates. Standard errors are shown in parentheses.

it entails extracting a confidence score for each LLM answer that we obtain and comparing it to the empirical hallucination rate that we observe. Table 7 reports the results of this analysis at the task level, and Figure 9 pools our findings at the LLM level by plotting those two metrics—confidences and empirical non-hallucination frequencies—against each other, binned into 10 equally-sized bins (represented by the dots). In a perfectly calibrated model, the confidences and empirical frequencies would be perfectly correlated along the $y = x$ diagonal.

Overall, we note that PaLM 2 (pooled ECE = 0.057), GPT 3.5 (pooled ECE = 0.099), and GPT 4 (pooled ECE = 0.190) are significantly better calibrated than Llama 2 (pooled ECE = 0.421). Interestingly, although GPT 4 is our best performing model in terms of raw hallucination rates (Figure 6), it is actually less calibrated than PaLM 2 and GPT 3.5, which are otherwise inferior. This suggests that even the newest and most advanced LLMs may not always be superior in all desirable senses—although GPT 4 is currently the LLM least prone to hallucination, our results imply that when it *does* hallucinate, it does so in a way that is more likely to mislead users than GPT 3.5 or PaLM 2.

Diving into the task-level results (Table 7), we see that across all LLMs, calibration is poorer on our more complex tasks, like **Doctrinal agreement**, and on tasks directed toward lower levels of the judicial hierarchy. ECE is also higher on our partially open-ended tasks such as **Court** and **Author**. In these tasks, the LLM has a large but finite universe of responses, and the high ECE for these tasks reflects the LLMs' tendencies to over-report on the most prominent or widely known members of the response set.

In all cases, the calibration error is in the positive direction: our LLMs systematically *overestimate* their confidence relative to their actual rate of hallucination.¹² This finding, too, suggests that users should exercise caution when interpreting LLMs' responses to legal queries, especially those of Llama 2. Not only may they receive a hallucinated response, but they may receive one that the LLM is overconfident in and liable to repeat again.

6 Discussion

We began this article with a question that has surged in salience over the last twelve months: Will AI systems like ChatGPT soon reshape the practice of law and democratize access to justice? Although there is much enthusiasm for LLMs' potential to revolutionize these domains, we highlight the problem




Figure 8. Hallucination rates by LLM, all contra-factual tasks pooled. Llama 2 is very unlikely to hallucinate on these tasks because it almost always rejects the premise in the question. However, this tendency also leads it to perform more poorly on tasks with correct premises (cf. Figure 6).

¹² In the [Online Appendix](#), we explore whether this bias can be corrected with an *ex post* scaling adjustment, but conclude that challenges remain.

Table 7 Expected calibration error (ECE) across levels of the federal judiciary

Task	Prompt	SCOTUS (1794–2015; n = 100)					USCOA (1895–2019; n = 100)					USDC (1932–2019; n = 100)				
		GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2
Existence	Zero-shot	0.026 (0.005)	0.008 (0.001)	0.119 (0.004)	0.262 (0.006)	0.124 (0.005)	0.007 (0.001)	0.182 (0.003)	0.117 (0.005)	0.074 (0.004)	0.004 (0.000)	0.183 (0.003)	0.178 (0.006)			
	Few-shot	0.149 (0.005)	0.038 (0.002)	0.041 (0.002)	0.998 (0.000)	0.073 (0.005)	0.064 (0.002)	0.014 (0.001)	0.998 (0.000)	0.024 (0.003)	0.053 (0.002)	0.038 (0.001)	0.096 (0.000)			
	Zero-shot	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.003 (0.000)	0.362 (0.006)	0.322 (0.006)	0.184 (0.006)	0.387 (0.006)	0.133 (0.005)	0.163 (0.005)	0.132 (0.005)	0.163 (0.006)			
Court	Few-shot	0 (0.000)	0.001 (0.000)	0 (0.000)	0.009 (0.001)	0.345 (0.006)	0.355 (0.006)	0.161 (0.006)	0.452 (0.008)	0.145 (0.005)	0.099 (0.004)	0.094 (0.005)	0.099 (0.005)			
	Zero-shot	0.113 (0.005)	0.069 (0.004)	0.026 (0.003)	0.068 (0.003)	0.143 (0.004)	0.073 (0.005)	0.004 (0.001)	0.036 (0.001)	0.104 (0.003)	0.043 (0.003)	0.002 (0.000)	0.022 (0.001)			
	Few-shot	0.191 (0.005)	0.029 (0.003)	0.059 (0.004)	0.063 (0.003)	0.138 (0.005)	0.036 (0.005)	0.01 (0.001)	0.032 (0.001)	0.099 (0.004)	0.018 (0.002)	0.001 (0.000)	0.051 (0.002)			
Author	Zero-shot	0.443 (0.006)	0.208 (0.006)	0.130 (0.005)	0.419 (0.006)	0.349 (0.005)	0.094 (0.003)	0.131 (0.002)	0.583 (0.004)	0.302 (0.005)	0.228 (0.005)	0.063 (0.002)	0.305 (0.004)			
	Few-shot	0.454 (0.006)	0.347 (0.006)	0.142 (0.005)	0.454 (0.005)	0.357 (0.004)	0.142 (0.003)	0.118 (0.003)	0.656 (0.004)	0.32 (0.004)	0.096 (0.003)	0.045 (0.002)	0.481 (0.005)			
	Zero-shot	0.203 (0.007)	0.431 (0.007)	0.291 (0.008)	0.199 (0.008)	0.17 (0.007)	0.557 (0.006)	0.382 (0.006)	0.148 (0.006)	0.099 (0.006)	0.004 (0.006)	0.002 (0.006)	0.005 (0.006)			
Disposition	Few-shot	0.239 (0.006)	0.314 (0.008)	0.165 (0.007)	0.283 (0.008)	0.165 (0.007)	0.205 (0.007)	0.227 (0.007)	0.439 (0.007)	0.32 (0.007)	0.096 (0.007)	0.045 (0.007)	0.481 (0.007)			
	Zero-shot	0.308 (0.025)	0.246 (0.022)	0.116 (0.019)	0.510 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)			
	Few-shot	0.377 (0.025)	0.68 (0.022)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)			
Overruling year	Zero-shot	0.308 (0.025)	0.246 (0.022)	0.116 (0.019)	0.510 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)			
	Few-shot	0.377 (0.025)	0.68 (0.022)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)			
	Zero-shot	0.308 (0.025)	0.246 (0.022)	0.116 (0.019)	0.510 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)	0.154 (0.019)	0.754 (0.019)			

Table 7. Continued

Task	Prompt	SCOTUS (1794–2015; n = 100)				USCOA (1895–2019; n = 100)				USDC (1932–2019; n = 100)			
		GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2	GPT 4	GPT 3.5	PaLM 2	Llama 2
Doctrinal agreement	Zero-shot	0.369 (0.006)	0.527 (0.007)	0.165 (0.006)	0.564 (0.006)								
	Few-shot	0.319 (0.007)	0.409 (0.008)	0.152 (0.007)	0.548 (0.007)								

^a 1810–2022 (n = 279)

^b 1796–2005 (n = 5,000)

Note: Table reports expected calibration error between empirical hallucination rates and estimated conditional probabilities. Conditional probabilities are estimated by sampling 10 responses from the model at temperature 1 and assessing their agreement with the model's greedy response. Bootstrapped standard errors are shown in parentheses.

of legal hallucinations, which remains a serious obstacle to the adoption of these models. Performing the first systematic empirical test of popular perceptions (Roberts 2023; Romoser 2023; Weiser 2023), we show that factual legal hallucinations are widespread in the LLMs that we study—OpenAI’s ChatGPT 4, OpenAI’s ChatGPT 3.5, Google’s PaLM 2, and Meta’s Llama 2—on the bulk of the legal knowledge tasks that we profile (Section 5.1).

We also push beyond conventional wisdom by documenting the correlates of these hallucinations and by surfacing two additional behaviors that threaten LLMs’ utility for legal applications: (i) their susceptibility to contra-factual bias, i.e., their inability to handle queries containing an erroneous or mistaken starting point (Section 5.2), and (ii) their certainty in their responses, i.e., their inability to always “know what they know” (Section 5.3). Unfortunately, we find that LLMs frequently provide seemingly genuine answers to legal questions whose premises are false by construction, and that under their default configurations they are imperfect predictors of their own tendency to confidently hallucinate legal falsehoods.

These findings complicate the existing literature that suggests that LLMs are performing increasingly well on a number of legal benchmarking tasks (Ash et al. 2024; Blair-Stanek, Holzenberger, and Van Durme 2023; Choi et al. 2022; Fei et al. 2023; Guha et al. 2023; Nay et al. 2023; Katz et al. 2023; Trozze, Davies, and Kleinberg 2023). Our study is related to this prior research, but is oriented in a slightly different vein. Instead of examining LLMs’ ability to engage in legal reasoning, we assess LLMs’ capacity to internalize legal knowledge. Ultimately, LLMs will need to excel in both of these respects if they are going to be effectively integrated into the legal profession. So long as they suffer from gaps in their background legal knowledge—as our results suggest—they will be unable to function as reliable sources of legal counsel and advice, no matter how strong their in-context reasoning abilities become.

Our results therefore temper optimism for the ability of off-the-shelf, publicly available LLMs to accelerate access to justice (Perlman 2023; Tan, Westermann, and Benykhlef 2023; Tito 2017). Indeed, our findings suggest that the risks of using these generic foundation models are especially high for litigants who are:

1. Filing in courts lower in the judicial hierarchy or those located in less prominent jurisdictions,
2. Seeking more complex forms of legal information,
3. Formulating questions with mistaken premises, or
4. Unsure of how much to trust the LLMs’ responses.

In short, we find that the risks are highest for those who would benefit from LLMs most—under-resourced or *pro se* litigants. Some of these risks—namely, (3) and (4)—might be mitigated with improved user




Figure 9. Calibration curves by LLM, all reference-based tasks pooled. PaLM 2 is best calibrated on legal queries, followed by GPT 3.5, GPT 4, and lastly Llama 2, which is significantly worse than the first three models.

education, but others—(1) and (2)—are more intractable. LLMs would ideally do best at localized legal information (rather than SCOTUS-level information), be able to correct users when they ask misguided questions (rather than accepting their premises at face value), and be able to moderate their responses with the appropriate level of confidence (rather than hallucinating with conviction). Consequently, we echo concerns that the proliferation of LLMs may ultimately exacerbate, rather than eradicate, existing inequalities in access to legal services (Draper and Gillibrand 2023; Simshaw 2022). At the same time, increased reliance on LLMs also has the potential to produce a kind of legal monoculture (Creel and Hellman 2022; Kleinberg and Raghavan 2021), with users being fed information from only a limited subset of judicial sources that elide many of the deeper nuances of the law. This new monoculture, in turn, is likely to reify the geographic, temporal, and judge-level biases that we diagnose above, as the foundation-like property of these models permits those biases to propagate into any downstream tools built on top of the original LLM (Bommasani et al. 2022).

Some recent research suggests that hallucinations can be diminished through the adoption of techniques like retrieval-augmented generation (RAG) (Shuster et al. 2021; Cui et al. 2023; Peng et al. 2023; Savelka et al. 2023), advanced prompting (such as chain-of-thought prompting or chain-of-verification) (Si et al. 2023; Lei et al. 2023; Mündler et al. 2023; Ji, Yu, et al. 2023; Dhuliwala et al. 2023; Suzgun and Kalai 2024), specialized fine-tuning (Tian, Mitchell, Yao, et al. 2023; Razumovskaia et al. 2023; Zhang et al. 2023), factuality-focused decoding methods (Shi et al. 2023; Mallen et al. 2023; Li et al. 2024; Chuang et al. 2024), or external database checks (Chern et al. 2023; Peng et al. 2023; Qin et al. 2023; Gou et al. 2024). These methods have shown promising results in significantly reducing hallucinated content and enhancing the accuracy, reliability, and faithfulness of model outputs. However, we caution that these approaches are not without limitations.

For example, the effectiveness of RAG-based methods heavily relies on the quality of their retrieval mechanisms (Wu et al. 2024). Moreover, accurately parsing and understanding the content of input queries poses a challenge, especially when queries are inherently ambiguous or irrelevant to the domain of focus (Tonmoy et al. 2024). Additionally, the task of retrieving relevant and precise information from extensive corpora can be computationally demanding and resource-intensive, necessitating continuous updating and modification of knowledge databases to keep pace with the latest information (Chen et al. 2023; Siriwardhana et al. 2023; Ram et al. 2023; Cheng et al. 2024). There may also be situations where the knowledge database might contain conflicting or contradictory information, making it unclear which pieces of relevant information to extract (Wang et al. 2023; Yu et al. 2023; Gao et al. 2023). For instance, when a legal case is overruled, or when there is a circuit split on an issue, the retrieval module must have some mechanism to distinguish outdated or jurisdictionally irrelevant sources from apposite and binding law.

Furthermore, methods for detecting hallucinations and evaluating their mitigation are themselves not foolproof. Evaluation datasets and metrics may not always accurately reflect real-world performance reliability (Ji et al. 2023; Lucas et al. 2023; Zhang et al. 2023). Biases could be embedded within the evaluation dataset, or the automated metric employed to quantify hallucination may lack comprehensiveness or task-specificity (Kang, Blevins, and Zettlemoyer 2024). Therefore, it is far from clear whether these technical improvements will be able to truly solve the hallucination problem.

Finally, we also emphasize that the challenges presented by legal hallucinations are not only empirical, but also normative. Although data-rich and moneyed players certainly stand at an advantage when it comes to building hallucination-free legal LLMs for their own private use, it is not clear that even infinite resources can entirely solve the conceptual problems we diagnose. As we discuss in Section 2.3, model fidelity to the training corpus, model fidelity to the user's prompt, and model fidelity to the facts of the world—i.e., the law—are normative commitments that stand in tension with each other, despite all being independently desirable technical properties of an LLM. Ultimately, since hallucinations of some kind are generally inevitable at the token level (Kalai and Vempala 2023; Xu, Jain, and Kankanhalli 2024), developers of legal LLMs will need to make choices about which type(s) of hallucinations to minimize, and they should make these choices transparent to their downstream users. Only then can individual litigants decide for themselves whether the legal information they seek to obtain from LLMs is trustworthy or not.

To aid in future research in this area, we release a test dataset of our queries and answers on the HuggingFace platform, which scholars can use to continue to evaluate LLMs as they advance in legal

sophistication.¹³ In the meantime, more experienced legal practitioners may find some value in consulting LLMs for certain tasks, but even these users should remain vigilant in their use, taking care to verify the accuracy of their prompts and the quality of their chosen LLM's responses. Similarly, legal scholars and educators seeking to use LLMs as automated research assistants (Livermore, Herron, and Rockmore 2024) or student aids (Choi and Schwarcz 2024) must be cautious to not inadvertently inject these LLMs' subtle knowledge biases into their own applications. Put differently, our findings underscore the importance of human-centered AI. Responsible integration of LLMs into legal tasks must augment lawyers, researchers, and litigants and not, as Chief Justice Roberts has put it, risk "dehumanizing the law" (Roberts 2023, 5).

Supplementary Material

Supplementary material is available at *JLA* online.

References

- Agrawal, Ayush, et al. 2023. Do Language Models Know When They're Hallucinating References? *Findings of the Association for Computational Linguistics: EACL 2024*, 912-928. St. Julian's, Malta: Association for Computational Linguistics. <https://aclanthology.org/2024.findings-eacl.62>
- Ambrogio, Bob. 2023. As Allen & Overy Deploys GPT-based Legal App Harvey Firmwide, Founders Say Other Firms Will Soon Follow. *LawSites*. <https://www.lawnext.com/2023/02/as-allen-overly-deploys-gpt-based-legal-app-harvey-firmwide-founders-say-other-firms-will-soon-follow.html>.
- Anil, Rohan, et al. 2023. PaLM 2 Technical Report. doi:10.48550/arXiv.2305.10403
- Ash, Elliott, et al. 2024. Translating Legalese: Enhancing Public Understanding of Court Opinions with Legal Summarizers In Proceedings of the Symposium on Computer Science and Law, 136-57. CSLAW '24. NY, USA: Association for Computing Machinery.
- Azaria, Amos, & Tom Mitchell. 2023. The Internal State of an LLM Knows When It's Lying. doi:10.48550/arXiv.2304.13734
- Bar-Gill, Oren, Cass R. Sunstein & Inbal Talgam-Cohen. 2023. Algorithmic Harm in Consumer Markets. 15 *J. Legal Anal.* 1-47.
- Barocas, Solon & Andrew D. Selbst. 2016. Big Data's Disparate Impact. 104 *California Law Rev.* 671-732.
- Ben-Shahar, Omri. 2023. Privacy Protection, At What Cost? Exploring the Regulatory Resistance to Data Technology in Auto Insurance. 15 *J. Legal Anal.* 129-57.
- Black, Ryan C & James F. Spriggs II. 2013. The Citation and Depreciation of U.S. Supreme Court Precedent. 10 *J. Empirical Legal Stud.* 325-58.
- Blair-Stanek, Andrew, Nils Holzenberger, and Benjamin Van Durme. 2023. "Can GPT-3 Perform Statutory Reasoning?" In Proceedings of the Nineteenth International Conference on Artificial Intelligence and Law. Braga, Portugal: Association for Computing Machinery.
- Bommasani, Rishi, et al. 2022. On the Opportunities and Risks of Foundation Models. doi:10.48550/arXiv.2108.07258
- Brown, Tom B., et al. 2020. Language Models Are Few-Shot Learners. doi: 10.48550/arXiv.2005.14165
- Cao, Meng, Yue Dong, & Jackie Cheung. 2022. Hallucinated but Factual! Inspecting the Factuality of Hallucinations in Abstractive Summarization. In Smaranda Muresan, Preslav Nakov, & Aline Villavicencio, eds., Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 3340-54. Dublin, Ireland: Association for Computational Linguistics.
- Cao, Ziqiang, et al. 2018. Faithful to the Original: Fact Aware Neural Abstractive Summarization. 32 *Proceedings of the AAAI Conference on Artificial Intelligence* 4784-4791.
- Caselaw Access Project. 2023. Caselaw Access Project. <https://case.law>.
- Carlini, Nicholas, et al. 2022. Quantifying Memorization Across Neural Language Models. doi:10.48550/arXiv.2202.07646
- Chann, Seherman. 2023. Non-Determinism in GPT-4 Is Caused by Sparse MoE. <https://152334H.github.io/blog/non-determinism-in-gpt-4/>.

¹³ https://huggingface.co/datasets/reglab/legal_hallucinations

- Chen, Jiawei, et al. 2023. Benchmarking Large Language Models in Retrieval-Augmented Generation. 38 AAAI Conference on Artificial Intelligence 17754–62. doi: [10.48550/arXiv.2309.01431](https://doi.org/10.48550/arXiv.2309.01431)
- Cheng, Xin, et al. 2024. Lift Yourself Up: Retrieval-augmented Text Generation with Self Memory. *37 Conference on Neural Information Processing Systems* 1–20.
- Chern, I, et al. 2023. FacTool: Factuality Detection in Generative AI—A Tool Augmented Framework for Multi-Task and Multi-Domain Scenarios. doi:[10.48550/arXiv.2307.13528](https://doi.org/10.48550/arXiv.2307.13528)
- Chien, Colleen V, et al. 2024. How Generative AI Can Help Address the Access to Justice Gap Through the Courts. *Loyola Los Angeles Law Rev.*, forthcoming.
- Choi, Jonathan H, et al. 2022. “ChatGPT Goes to Law School. 71 *J. Legal Educ.* 387–400.
- Choi, Jonathan H & Daniel Schwarcz. 2024. AI Assistance in Legal Analysis: An Empirical Study. *J. Legal Educ.* doi: [10.2139/ssrn.4539836](https://doi.org/10.2139/ssrn.4539836). forthcoming
- Chuang, Yung-Sung, et al. 2024. “DoLa: Decoding by Contrasting Layers Improves Factuality in Large Language Models.” In *The Twelfth International Conference on Learning Representations*.
- Congress.gov. 2023. Table of Supreme Court Decisions Overruled by Subsequent Decisions. <https://constitution.congress.gov/resources/decisions-overruled/>.
- Creel, Kathleen & Deborah Hellman. 2022. The Algorithmic Leviathan: Arbitrariness, Fairness, and Opportunity in Algorithmic Decision-Making Systems. 52 *Canad. J. Philos.* 26–43.
- Cui, Jiayi, et al. 2023. ChatLaw: Open-Source Legal Large Language Model with Integrated External Knowledge Bases. doi:[10.48550/arXiv.2306.16092](https://doi.org/10.48550/arXiv.2306.16092)
- Cushman, Jack, Matthew Dahl & Michael Lissner. 2021. Eyecite: A Tool for Parsing Legal Citations. 6 *Open Source Software* 3617.
- Dell’Acqua, Fabrizio, et al. 2023. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. doi:[10.2139/ssrn.4573321](https://doi.org/10.2139/ssrn.4573321)
- Deroy, Aniket, Kripabandhu Ghosh, and Saptarshi Ghosh. 2023. How Ready Are Pre-trained Abstractive Models and LLMs for Legal Case Judgement Summarization? doi:[10.48550/arXiv.2306.01248](https://doi.org/10.48550/arXiv.2306.01248)
- Dhuliawala, Shehzaad, et al. 2023. Chain-of-Verification Reduces Hallucination in Large Language Models. *arXiv preprint arXiv:2309.11495* doi: [10.48550/arXiv.2309.11495](https://doi.org/10.48550/arXiv.2309.11495)
- Draper, Chris, and Nicky Gillibrand. 2023. The Potential for Jurisdictional Challenges to AI or LLM Training Datasets. In *Proceedings of the ICAIL 2023 Workshop on Artificial Intelligence for Access to Justice*. Braga, Portugal: CEUR Workshop Proceedings.
- Dworkin, Ronald. 1986. *Law’s Empire*. Cambridge, MA: Harvard University Press.
- Eloundou, Tyna, et al. 2023. GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. doi:[10.48550/arXiv.2303.10130](https://doi.org/10.48550/arXiv.2303.10130)
- Engel, Christoph & Nina Grgić-Hlačić. 2021. Machine Advice with a Warning about Machine Limitations: Experimentally Testing the Solution Mandated by the Wisconsin Supreme Court. 13 *J. Legal Anal.* 284–340.
- Engstrom, David Freeman, Daniel E. Ho. 2020. Algorithmic Accountability in the Administrative State. 37 *Yale J. Regul.* 800–54.
- Engstrom, David Freeman, et al. 2020. Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies. Technical report. Administrative Conference of the United States.
- Fei, Zhiwei, et al. 2023. LawBench: Benchmarking Legal Knowledge of Large Language Models. doi:[10.48550/arXiv.2309.16289](https://doi.org/10.48550/arXiv.2309.16289)
- Feijo, Diego de Vargas, & Viviane P. Moreira. 2023. Improving Abstractive Summarization of Legal Rulings through Textual Entailment. 31 *Artif. Intell. Law* 91–113.
- Feng, Shangbin, et al. 2024. Don’t Hallucinate, Abstain: Identifying LLM Knowledge Gaps via Multi-LLM Collaboration. doi:[10.48550/arXiv.2402.00367](https://doi.org/10.48550/arXiv.2402.00367)
- Fowler, James H, et al. 2007. Network Analysis and the Law: Measuring the Legal Importance of Precedents at the U.S. Supreme Court. 15 *Polit. Anal.* 324–46.
- Gao, Yunfan, et al. 2023. Retrieval-Augmented Generation for Large Language Models: A Survey. doi:[10.48550/arXiv.2312.10997](https://doi.org/10.48550/arXiv.2312.10997)
- Gillis, Talia B & Jann L. Spiess. 2019. Big Data and Discrimination. 86 *Univ. Chicago Law Rev.* 459–88.
- Gou, Zhibin, et al. 2024. CRITIC: Large Language Models Can Self-Correct with Tool-Interactive Critiquing. In *The Twelfth International Conference on Learning Representations*.
- Guha, Neel, et al. 2023. LegalBench: A Collaboratively Built Benchmark for Measuring Legal Reasoning in Large Language Models. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

- Henderson, Peter, Tatsunori Hashimoto & Mark Lemley. 2023. Where's the Liability in Harmful AI Speech?. *3 J Free Speech Law* 589–650.
- Henderson, Peter, et al. 2022. Pile of Law: Learning Responsible Data Filtering from the Law and a 256GB Open-Source Legal Dataset. doi:[10.48550/arXiv.2207.00220](https://doi.org/10.48550/arXiv.2207.00220).
- Ji, Ziwei, et al. 2023. Survey of Hallucination in Natural Language Generation. *55 ACM Comput. Surv.* 248:1–248:38.
- Ji, Ziwei, et al. 2023. Towards Mitigating LLM Hallucination via Self Reflection. In Houda Bouamor, Juan Pino, & Kalika Bali, eds., *Findings of the Association for Computational Linguistics: EMNLP 2023*, 1827–43. Singapore: Association for Computational Linguistics.
- Jones, Erik, and Jacob Steinhardt. 2022. Capturing Failures of Large Language Models via Human Cognitive Biases. *36 Conference on Neural Information Processing Systems* 1–15. doi:[10.48550/arXiv.2202.1229](https://doi.org/10.48550/arXiv.2202.1229).
- Kadavath, Saurav, et al. 2022. Language Models (Mostly) Know What They Know. doi:[10.48550/arXiv.2207.05221](https://doi.org/10.48550/arXiv.2207.05221)
- Kalai, Adam Tauman, & Santosh S. Vempala. 2023. Calibrated Language Models Must Hallucinate. doi:[10.48550/arXiv.2311.14648](https://doi.org/10.48550/arXiv.2311.14648)
- Kang, Haoqiang, Terra Blevins, & Luke Zettlemoyer. 2024. Comparing Hallucination Detection Metrics for Multilingual Generation. doi:[10.48550/arXiv.2402.10496](https://doi.org/10.48550/arXiv.2402.10496)
- Katz, Daniel Martin, et al. 2023. GPT-4 Passes the Bar Exam. doi:[10.2139/ssrn.4389233](https://doi.org/10.2139/ssrn.4389233)
- King, Jennifer, et al. 2023. The Privacy- Bias Tradeoff: Data Minimization and Racial Disparity Assessments in U.S. Government. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, FAccT '23*, 492–505. Chicago, IL, USA: Association for Computing Machinery.
- Kleinberg, Jon, et al. 2018. Discrimination in the Age of Algorithms. *10 J. Legal Anal.* 113–74.
- Kleinberg, Jon & Manish Raghavan. 2021. Algorithmic Monoculture and Social Welfare. *118 Proc. Natl. Acad. Sci. U.S.A* 1–7.
- Krishna, Kalpesh, Aurko Roy, & Mohit Iyyer. 2021. Hurdles to Progress in Long-form Question Answering. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 4940–57. Association for Computational Linguistics.
- Kuersten, Ashlyn K., & Susan B. Haire. 2011. Update to the Appeals Courts Database (1997–2002). <https://artsandsciences.sc.edu/poli/juri/appct.htm>
- Lee, Minhyeok. 2023. A Mathematical Investigation of Hallucination and Creativity in GPT Models. *11 Mathematics* 2320.
- Lee, Nayeon, et al. 2023. Factuality Enhanced Language Models for Open-Ended Text Generation. doi:[10.48550/arXiv.2206.04624](https://doi.org/10.48550/arXiv.2206.04624)
- Lei, Deren, et al. 2023. Chain of Natural Language Inference for Reducing Large Language Model Ungrounded Hallucinations. doi:[10.48550/arXiv.2310.03951](https://doi.org/10.48550/arXiv.2310.03951)
- Lemley, Mark A & Bryan Casey. 2019. Remedies for Robots. *86 Univ. Chicago L. Rev.* 1311–96.
- Li, Junyi, et al. 2023. HaluEval: A Large-Scale Hallucination Evaluation Benchmark for Large Language Models. doi: [10.48550/arXiv.2305.11747](https://doi.org/10.48550/arXiv.2305.11747)
- Li, Kenneth, et al. 2024. Inference-Time Intervention: Eliciting Truthful Answers From a Language Model.
- Lin, Stephanie, Jacob Hilton, and Owain Evans. 2022. TruthfulQA: Measuring How Models Mimic Human Falsehoods. doi:[10.48550/arXiv.2109.07958](https://doi.org/10.48550/arXiv.2109.07958)
- Livermore, Michael A, Felix Herron & Daniel Rockmore. 2024. Language Model Interpretability and Empirical Legal Studies. *J. Institut. Theoret. Econ.*, forthcoming
- Lucas, Jason, et al. 2023. Fighting Fire with Fire: The Dual Role of LLMs in Crafting and Detecting Elusive Disinformation. In Houda Bouamor, Juan Pino, and Kalika Bali eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 14279–305. Singapore: Association for Computational Linguistics.
- Luft, Joseph, and Harrington Ingham. 1955. The Johari Window as a Graphic Model of Interpersonal Awareness. In *Proceedings of the Western Training Laboratory in Group Development*. University of California, Los Angeles, Extension Office.
- Mallen, Alex, et al. 2023. When Not to Trust Language Models: Investigating Effectiveness of Parametric and Non-Parametric Memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 9802–22.
- Manakul, Potsawee, Adian Liusie, & Mark J. F. Gales. 2023. SelfCheckGPT: Zero-Resource Black-Box Hallucination Detection for Generative Large Language Models. doi:[10.48550/arXiv.2303.08896](https://doi.org/10.48550/arXiv.2303.08896)

- Maynez, Joshua, et al. 2020. On Faithfulness and Factuality in Abstractive Summarization. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 1906–19. Association for Computational Linguistics.
- Mayson, Sandra G. 2019. Bias In, Bias Out. 128 *Yale Law J.* 2122–473.
- Min, Sewon, et al. 2023. FACTScore: Fine-grained Atomic Evaluation of Factual Precision in Long Form Text Generation. doi:10.48550/arXiv.2305.14251
- Mündler, Niels, et al. 2023. Self-Contradictory Hallucinations of Large Language Models: Evaluation, Detection and Mitigation. doi: 10.48550/arXiv.2305.15852
- Nay, John J., et al. 2023. Large Language Models as Tax Attorneys: A Case Study in Legal Capabilities Emergence. doi:10.48550/arXiv.2306.07075
- OpenAI. 2023a. GPT-4 Technical Report. doi:10.48550/arXiv.2303.08774
- . 2023b. *Introducing ChatGPT*. <https://openai.com/blog/chatgpt>
- Peng, Baolin, et al. 2023. Check Your Facts and Try Again: Improving Large Language Models with External Knowledge and Automated Feedback. doi:10.48550/arXiv.2302.12813.
- Perlman, Andrew. 2023. The Implications of ChatGPT for Legal Services and Society. *The Practice*, March/April. <https://clp.law.harvard.edu/knowledge-hub/magazine/issues/generative-ai-in-the-legal-profession/the-implications-of-chatgpt-for-legal-services-and-society/>.
- Qin, Yujia, et al. 2023. Tool Learning With Foundation Models. doi:10.48550/arXiv.2304.08354
- Ram, Ori, et al. 2023. In-Context Retrieval-Augmented Language Models. 11 *Trans. Assoc. Computat. Linguist.* 1316–31.
- Razumovskaia, Evgeniia. 2023. Dial BeInfo for Faithfulness: Improving Factuality of Information-Seeking Dialogue via Behavioural Fine-Tuning. doi:10.48550/arXiv.2311.09800
- Roberts, John G. 2023. 2023 Year-End Report on the Federal Judiciary. <https://www.supremecourt.gov/publicinfo/year-end/2023year-endreport.pdf>
- Rodgers, Ian, John Armour & Mari Sako. 2023. How Technology Is (or Is Not) Transforming Law Firms. 19 *Ann. R. Law Social Sci.* 299–317.
- Romoser, James. 2023. No, Ruth Bader Ginsburg Did Not Dissent in Obergefell — and Other Things ChatGPT Gets Wrong about the Supreme Court. *SCOTUSblog*. <https://www.scotusblog.com/2023/01/no-ruth-bader-ginsburg-did-not-dissent-in-obergefell-and-other-things-chatgpt-gets-wrong-about-the-supreme-court/>.
- Savelka, Jaromir, et al. 2023. Explaining Legal Concepts with Augmented Large Language Models (GPT-4). doi:10.48550/arXiv.2306.09525
- Sharma, Mrinank, et al. 2023. Towards Understanding Sycophancy in Language Models. doi:10.48550/arXiv.2310.13548
- Shi, Weijia, et al. 2023. Trusting Your Evidence: Hallucinate Less with Context-aware Decoding. doi:10.48550/arXiv.2305.14739
- Shuster, Kurt, et al. 2021. “Retrieval Augmentation Reduces Hallucination in Conversation.” In Findings of the Association for Computational Linguistics: EMNLP 2021, 3784–803. Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Si, Chenglei, et al. 2023. “Prompting GPT-3 To Be Reliable.” In *The Eleventh International Conference on Learning Representations*.
- Simshaw, Drew. 2022. Access to A.I. Justice: Avoiding an Inequitable Two-Tiered System of Legal Services. 24 *Yale J. Law Technol.* 150–226.
- Siriwardhana, Shamane, et al. 2023. Improving the Domain Adaptation of Retrieval Augmented Generation (RAG) Models for Open Domain Question Answering. 11 *Trans. Assoc. Comput. Linguist.* 1–17.
- Solow-Niederman, Alicia. 2020. Administering Artificial Intelligence. 93 *Southern California Law Rev.* 633–96.
- Songer, Donald R. 2008. *The United States Courts of Appeals Database, 1925–1996*.
- Spaeth, Harold J., et al. 2022. *2022 Supreme Court Database, Version 2022 Release 01*. <http://supremecourtdatabase.org/>
- Suri, Gaurav, et al. 2023. Do Large Language Models Show Decision Heuristics Similar to Humans? A Case Study Using GPT-3.5. doi: 10.48550/arXiv.2305.04400
- Suzgun, Mirac, & Adam Tauman Kalai. 2024. *Meta-prompting: Enhancing Language Models with Task-Agnostic Scaffolding*. doi:10.48550/arXiv.2401.12954
- Tan, Jinzhe, Hannes Westermann, & Karim Benyekhlef. 2023. “ChatGPT as an Artificial Lawyer?” In Proceedings of the ICAIL 2023 Workshop on Artificial Intelligence for Access to Justice. Braga, Portugal: CEUR Workshop Proceedings.

- Tian, Katherine, et al. 2023. *Fine-Tuning Language Models for Factuality*. doi:10.48550/arXiv.2311.08401
- Tian, Katherine, et al. 2023. *Just Ask for Calibration: Strategies for Eliciting Calibrated Confidence Scores from Language Models Fine-Tuned with Human Feedback*. doi:10.48550/arXiv.2305.14975
- Tito, Joel. 2017. *How AI Can Improve Access to Justice*. <https://www.centreforpublicimpact.org/insights/joel-tito-ai-justice>
- Tonmoy, SM, et al. 2024. *A Comprehensive Survey of Hallucination Mitigation Techniques in Large Language Models*. doi:10.48550/arXiv.2401.01313
- Touvron, Hugo, et al. 2023. *Llama 2: Open Foundation and Fine-Tuned Chat Models*. doi:10.48550/arXiv.2307.09288
- Trozze, Arianna, Toby Davies, & Bennett Kleinberg. 2023. *Large Language Models in Cryptocurrency Securities Cases: Can ChatGPT Replace Lawyers?* doi:10.48550/arXiv.2308.06032
- Tversky, Amos & Daniel Kahneman. 1974. *Judgment under Uncertainty: Heuristics and Biases*. 185 *Science* 1124–31.
- van Deemter, Kees. 2024. *The Pitfalls of Defining Hallucination*. *Comput. Linguist.* 1–10, forthcoming
- Volokh, Eugene. 2023. *Large Label Models? Liability for AI Output*. 3 *J. Free Speech Law* 489–558.
- Wang, Cunxiang, et al. 2023. *Survey on Factuality in Large Language Models: Knowledge, Retrieval and Domain-Specificity*. doi:10.48550/arXiv.2310.07521
- Wei, Jerry, et al. 2023. *Simple Synthetic Data Reduces Sycophancy in Large Language Models*. doi:10.48550/arXiv.2308.03958
- Weiser, Benjamin. 2023. *Here's What Happens When Your Lawyer Uses ChatGPT*. *New York Times*. <https://www.nytimes.com/2023/05/27/nyregion/avianca-airline-lawsuit-chatgpt.html>.
- Wittgenstein, Ludwig. 1998 [1921]. *Tractatus Logico-Philosophicus*. Translated by C. K. Ogden. Dover.
- Wu, Kevin, et al. 2024. *How Well Do LLMs Cite Relevant Medical References? An Evaluation Framework and Analyses*. doi:10.48550/arXiv.2402.02008
- Xiong, Miao, et al. 2023. *Can LLMs Express Their Uncertainty? An Empirical Evaluation of Confidence Elicitation in LLMs*. doi:10.48550/arXiv.2306.13063
- Xu, Weijia, et al. 2023. *Understanding and Detecting Hallucinations in Neural Machine Translation via Model Introspection*. doi:10.48550/arXiv.2301.07779
- Xu, Ziwei, Sanjay Jain, & Mohan Kankanhalli. 2024. *Hallucination Is Inevitable: An Innate Limitation of Large Language Models*. doi:10.48550/arXiv.2401.11817
- Yin, Zhangyue, et al. 2023. *Do Large Language Models Know What They Don't Know?* doi:10.48550/arXiv.2305.18153
- Yu, Wenhao, et al. 2023. *Chain-of-Note: Enhancing Robustness in Retrieval-Augmented Language Models*. doi:10.48550/arXiv.2311.09210
- Zhang, Hanning, et al. 2023. *R-Tuning: Teaching Large Language Models to Refuse Unknown Questions*. doi:10.48550/arXiv.2311.09677
- Zhang, Yue, et al. 2023. *Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models*. doi:10.48550/arXiv.2309.01219.
- Zhang, Yunfeng, Q. Vera Liao, & Rachel K. E. Bellamy. 2020. *Effect of Confidence and Explanation on Accuracy and Trust Calibration in AI-assisted Decision Making*. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 295–305. FAT* '20. NY, USA: Association for Computing Machinery.

TRADEMARK SEARCH, ARTIFICIAL INTELLIGENCE, AND THE ROLE OF THE PRIVATE SECTOR

Sonia K. Katyal[†] & *Aniket Kesari*^{††}

ABSTRACT

Almost every industry today is confronting the potential role that artificial intelligence and machine learning can play in its future. While there are many, many studies on the role of AI in marketing to the consumer, there is less discussion of the role of AI in creating and selecting a trademark that is both distinctive, recognizable, and meaningful to the average consumer. As we argue, given that the role of AI is rapidly increasing in trademark search and similarity areas, lawyers and scholars should be apprised of some of the dramatic implications that AI's role can produce.

We begin, mainly, by proposing that AI should be of interest to anyone studying trademarks and the role that they play in economic decision-making. By running a series of empirical experiments regarding search, we show how comparative work can help us to assess the efficacy of various trademark search engines, many of which draw on a variety of machine learning methods. Traditional approaches to trademarks, spearheaded by economic approaches, have focused almost exclusively on consumer-based, demand-side considerations regarding search. Yet, as we show in this paper, these approaches are incomplete because they fail to take into account the substantial costs that are also faced by not just consumers, but trademark applicants as well. In the end, as we show, machine learning techniques will have a transformative effect on the application and interpretation of foundational trademark doctrines, producing significant implications for the trademark ecosystem. In an age where AI will increasingly govern the process of trademark selection, we argue that the classic division between consumers and trademark owners is perhaps deserving of an updated, supply-side framework. As we argue, a new framework is needed—one that reflects that putative trademark owners, too, are also consumers in the trademark selection ecosystem, and that this insight has transformative potential for encouraging both innovation and efficiency.

DOI: <https://doi.org/10.15779/Z380V89H87>

© 2020 Sonia K. Katyal & Aniket Kesari.

[†] Haas Distinguished Chair, University of California, Berkeley; Co-Associate Dean for Faculty Research and Development.

^{††} Postdoctoral Scholar, Social Sciences D-Lab, University of California at Berkeley; JD Candidate, Yale University. Thanks to Bob Cooter, Stacey Dogan, Aaron Edlin, Benoit Fallenius, Jeanne Fromer, Su Li, Trevor Little, Lisa Larrimore Ouellette, Frank Partnoy, Justin McCrary, Steven Davidoff Solomon, Adam Sterling, Jared Elias, James Hicks, Eric Goldman, Amelia Miazad, Tejas Narechania, Prasad Krishnamurthy, Steven Davidoff Solomon, Robert Bartlett, Tabrez Ebrahim, Barton Beebe, Simone Ross, and Glynn Lunney for their comments. This paper benefited from comments received at a conference with the Berkeley Center for Law, Commerce, and the Economy. We especially thank Jim Dempsey for his helpful insights on the project, the editors at *BIIJ* for their work, and Mehtab Khan for helpful research assistance. This research was supported by a grant from Markify. The authors affirm that they had complete freedom to direct the study, analysis, results, and reporting, without influence, editorial direction, or censorship.

© 2020 by Sonia K. Katyal and Aniket Kesari. Originally published in 35 *Berkeley Tech. L. J.* 501 (2020). Reprinted with kind permission of the authors.

TABLE OF CONTENTS

INTRODUCTION.....	503
I. SEARCH COSTS IN TRADEMARK LAW: A VIEW FROM THE CONSUMER.....	508
A. SEARCH, EXPERIENCE, AND CREDENCE ATTRIBUTES IN CONSUMER DECISION-MAKING.....	509
B. TRADEMARK LAW AND CONSUMER SEARCH COSTS	511
II. SEARCH COSTS IN TRADEMARK REGISTRATION: A VIEW FROM A TRADEMARK APPLICANT	514
A. SUPPLEMENTING TRADEMARK SEARCH IN THE PRIVATE SECTOR ..	516
B. ARTIFICIAL INTELLIGENCE AND THE ADMINISTRATION OF INTELLECTUAL PROPERTY	522
C. ARTIFICIAL INTELLIGENCE IN PRIVATE TRADEMARK SEARCH AND REGISTRATION.....	527
1. <i>Search, Identification, and Suggestion</i>	531
2. <i>Registration and Clearance</i>	531
3. <i>Comparison and Determining Substantial Similarity</i>	532
4. <i>Prediction and Risk Assessment</i>	533
5. <i>Brand Management</i>	533
III. A COMPARATIVE ASSESSMENT OF THE PRIVATE SECTOR IN TRADEMARK SEARCH.....	534
A. TRADEMARK SEARCH AND REGISTRATION PROCEDURES.....	538
B. A COMPARISON OF TRADEMARK SEARCH ENGINES.....	542
1. <i>Public Search Engines</i>	542
a) USPTO	542
2. <i>Private Search Engines</i>	545
a) Corsearch.....	545
b) Markify.....	547
c) Trademarkia	549
d) TrademarkNow	550
3. <i>Our Methodology</i>	551
4. <i>Generating Conflicted Trademarks</i>	552
5. <i>Scraping Websites</i>	554
6. <i>Exploratory Data Analysis</i>	555
a) Baseline.....	556
b) Exploratory Analysis of Private Search Engines.....	561
7. <i>Metrics</i>	567
8. <i>Results</i>	568
IV. IMPLICATIONS FOR FURTHER STUDY	572
A. OUTCOMES AND IMPLICATIONS	573
B. FRAMING TRADEMARK REGISTRATION AS AN ADVERSARIAL	

	MACHINE LEARNING PROBLEM.....	576
C.	RISK ASSESSMENT IN THE TRADEMARK ECOSYSTEM.....	578
D.	FUTURE WORK	587
CONCLUSION		588

INTRODUCTION

Almost every industry today is confronting the potential role that artificial intelligence (AI) and machine learning can play in its future. Intellectual Property (IP) and Information Law are no exception. In areas involving IP, many entities are studying the potential effect of descriptive and predictive analytics on its creation, registration, comparison, and litigation. The U.S. Patent and Trademark Office (USPTO) recently solicited public comments on the relationship between AI and IP,¹ held a conference on the subject, and even ran a contest for improving patent search with AI.² More recently, several prominent studies have focused on the role that machine learning can play at the USPTO in the process of prosecution.³

In the area of copyright law, scholars and commentators have voiced significant debate over whether AI-created works can be registered, and the role of human oversight in the crafting of authorship.⁴ There are fascinating

1. See *Request for Comments on Intellectual Property Protection for Artificial Intelligence Innovation*, FEDERAL REGISTER (Oct. 30, 2019), <https://www.federalregister.gov/documents/2019/10/30/2019-23638/request-for-comments-on-intellectual-property-protection-for-artificial-intelligence-innovation>; see also Neil Wilkof, *USPTO Conference on Artificial Intelligence and IP: A Report*, THE IPKAT (Mar. 20, 2019), <http://ipkitten.blogspot.com/2019/03/uspto-conference-on-artificial.html>.

2. See *USPTO's Challenge to Improve Patent Search with Artificial Intelligence*, GOVTRIBE (last updated Nov. 7, 2018), <https://govtribe.com/opportunity/federal-contract-opportunity/uspto-s-challenge-to-improve-patent-search-with-artificial-intelligence-rfusptoipatentsearch18>.

3. See generally Arti K. Rai, *Machine Learning at the Patent Office: Lessons for Patents and Administrative Law*, 104 IOWA L. REV. 2617 (2019). In this paper, we draw on Rai's instructive description of machine learning, which notes that "a distinctive feature of the genre is that the learning algorithm does not represent the decision rule; instead, the algorithm "learns" the decision rules from data known as training data." *Id.* (citing David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653 (2017) (explaining machine learning processes)); see also Tabrez Y. Ebrahim, *Automation & Predictive Analytics in Patent Prosecution: USPTO Implications & Policy*, 35 GA. ST. U.L. REV. 1185 (2019).

4. For a lengthier discussion of this literature and the relevant questions, see generally Jane C. Ginsburg & Luke Ali Budiardjo, *Authors and Machines*, 34 BERKELEY TECH. L.J. 343 (2019); Shyam Balganes, *Causing Copyright*, 117 COLUM. L. REV. 1 (2017).

questions about who owns the rights to an AI-generated work. Does the author of a program, the user, or the AI itself possess the intellectual property rights over these types of works? Determining the scope of authorship in an era where machines are increasingly capable of performing human-like tasks is a fascinating area of IP scholarship.⁵ Further, it promises to yield rich debates about the limits of property, personhood, and creativity.

Yet, surprisingly, very little legal scholarship has addressed the potential role for AI in the context of trademarks.⁶ For example, in December 2019, the World Intellectual Property Organization (WIPO) Secretariat issued a draft paper on IP and AI, and while it addressed a range of issues involving the administration of IP and other topics relating to patents, copyright, data, design, and capacity building, it did not cover trademarks.⁷ Similarly, while there are many studies on the role of AI in consumer marketing, there is very little scholarly research on the potential role of AI in the corresponding trademark ecosystem.⁸ This absence is surprising, especially considering that business owners continue to emphasize that trademarks are the most important area of IP protection.⁹ In the United States, IP-related industries

5. For a discussion of the intersection with trademark law and economics, see WORLD INTELLECTUAL PROP. ORG., 2013 WORLD INTELLECTUAL PROPERTY REPORT: BRAND – REPUTATION AND IMAGE IN THE GLOBAL MARKETPLACE, 81–108 (2013), https://www.wipo.int/edocs/pubdocs/en/wipo_pub_944_2013-chapter2.pdf.

6. There are very few law-related papers addressing trademarks and AI at the time of publication. See, e.g., Dev Gangjee, *Eye, Robot: Artificial Intelligence and Trade Mark Registers*, in TRANSITION AND COHERENCE IN INTELLECTUAL PROPERTY LAW (N. Bruun, G. Dinwoodie, M. Levin & A. Ohly eds., forthcoming 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3467627; Anke Moerland & Conrado Freitas, *Artificial Intelligence and Trade Mark Assessment*, in *Artificial Intelligence & Intellectual Property* (R. Hilty, K-C. Liu & J-A. Lee eds., forthcoming 2021), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3683807.

7. See WIPO Conversation on Intellectual Property (IP) and Artificial Intelligence (AI): Second Session, WIPO, https://www.wipo.int/meetings/en/details.jsp?meeting_id=55309 (last visited Jan. 22, 2021).

8. See, e.g., Thomas Davenport, Abhijit Guha, Dhruv Grewal & Timna Bressgott, *How Artificial Intelligence Will Change the Future of Marketing*, J. ACAD. MKTG. SCI. (2019), available for download at https://ideas.repec.org/a/spr/joamsc/v48y2020i1d10.1007_s11747-019-00696-0.html; Jan Keitzmann, Jeannette Paschen & Emily Treen, *Artificial Intelligence in Advertising: How Marketers Can Leverage Artificial Intelligence Along the Consumer Journey*, 58 J. ADVERT. RES. 263 (2018); Mónica Casabayó, Nuria Agell & Juan Carlos Aguado, *Using AI Techniques in the Grocery Industry: Identifying the Customers Most Likely to Defect*, 14 INT'L REV. RETAIL DISTRIB. & CONSUMER RES. 295 (2007); Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995 (2014) (offering a look into how technology-mediated advertising intersects with behavioral economics).

9. See *Trademarks, Copyright and Patents: Should Business Owners Really Care About IP?*, VARNUM (May 01, 2019), <https://www.varnumlaw.com/newsroom-publications-trademarks-copyrights-and-patents-why-business-owners-should-care-about-ip> (“A trademark is one of

support at least forty-five million U.S. jobs, contributing over thirty-eight percent to U.S. GDP.¹⁰

In this Article, we seek to remedy the absence of research in this field by studying the impact of AI on private trademark search engines and their economic and legal implications.¹¹ We begin by proposing, as a general matter, that AI should be of interest to anyone studying trademarks and the role that they play in economic decision-making. AI will fundamentally transform the trademark ecosystem, and the law will need to evolve as a result. The largest set of questions, we predict, emerges from the need for a more sophisticated approach regarding the impact of AI on the private sector of trademark search. As industries increasingly choose to rely on private AI-powered techniques for search, it becomes more and more essential to consider the nature of these technologies and their implications for trademark creation, comparison, and protection.

In turn, we argue that machine learning will have a transformative effect on the application and interpretation of foundational trademark doctrines. Our study focuses on the application of AI to trademark search and how it fits into a broader discussion about how AI will transform the economics of IP. Most traditional analyses of trademarks focus on the clarifying role of trademarks in aiding consumer search and demand for products in the marketplace. However, we believe that AI carries significant potential to affect the registration and quality of trademarks within the trademark ecosystem, thereby making it necessary to consider the effect of AI on trademark supply as well. Recent increases in trademark applications have exacerbated concerns regarding trademark quality; at least one study has observed, “. . . examiners are going through the motions to meet quota numbers and are not actually

the most important business assets that a company will ever own because it identifies and distinguishes the company and its products/services in the marketplace from its competitors.”); see also Darren Heitner, *Why Intellectual Property is Important for Your Business and What You Should be Doing Now to Protect It*, INC.COM (May 31, 2018), <https://www.inc.com/darren-heitner/why-intellectual-property-is-important-for-your-business-what-you-should-be-doing-now-to-protect-it.html> (discussing the importance of trademarks).

10. Robert Silvers, Sarah Pearce, Brad Newman, John Phillips, Elena Baca, Tom Brown, Scott Flicker, Emily Pidot, Carson Sullivan & Edward George, *Containing Risk and Seizing Opportunity: The In-house Lawyer's Guide to Artificial Intelligence*, PAUL HASTINGS LLP (Mar. 26, 2019), <https://www.paulhastings.com/publications-items/details/?id=43b9226d-2334-6428-811c-ff00004cbded>.

11. For a good discussion of various issues that have arisen in the recent rise of trademark applications, see *The Pressure of Rising Demand*, WORLD TRADEMARK REV. (July 1, 2016), <https://www.worldtrademarkreview.com/governmentpolicy/pressure-rising-demand> [hereinafter WTR Report] (noting rise in application filings and describing the role of the private sector).

examining the evidence.”¹² Thus, scholars are increasingly paying attention to the possibility that AI can, and should, be used by the government to even the playing field between itself and potential registrants, in order to improve the quality of registered IP.¹³ As AI tools proliferate in the private sector, government failure to adapt could exacerbate market inefficiencies stemming from information asymmetries.¹⁴

Since there are more trademarks than ever, searching them manually carries enormous costs. Private search algorithms reduce these costs by helping individuals traverse massive datasets efficiently, drawing on AI to do so. While a traditional trademark applicant might rely on government-supported techniques, the Trademark Electronic Search System (TESS), for searching confusingly similar marks, it turns out that TESS is often incomplete. Because of these gaps, several private trademark search engines have emerged to supplement TESS, using machine learning to provide more thorough results. However, not all AI-powered searches are created equal, and their efficacy is a key factor in determining whether users avoid the costs associated with a failed search. Each search engine uses its own methods, algorithms, and techniques to return results. These search engines generally aim to provide a user with a more comprehensive list of potential mark conflicts and to recommend whether the user should proceed with their trademark application, among other services.

As we argue in this Article, a high-level study of AI in the trademark search ecosystem offers us several contributions. To explore the intersection between TESS and private search engines, we conducted a series of experiments to compare the performance of AI-powered search engines in identifying potential conflicts under Section 2(d) of the Trademark Act, 15 U.S.C. § 1052(d),¹⁵ which forbids the registration of a trademark that is confusingly similar to an existing registered trademark. By running a series of comparisons

12. *See id.* at 3 (quoting a law firm in its survey responses).

13. *See, e.g.,* Ebrahim, *Automation & Predictive Analytics*, *supra* note 3, at 1188–89 (proposing that the magnified information asymmetries between the inventor and patent examiner can be reduced through artificial intelligence technology).

14. *See id.* at 1189, 1211–28.

15. 15 U.S.C. § 1052(d) (2018). The statute states:

No trademark by which the goods of the applicant may be distinguished from the goods of others shall be refused registration on the principal register on account of its nature unless it . . . [c]onsists of or comprises a mark which so resembles a mark registered in the Patent and Trademark Office, or a mark or trade name previously used in the United States by another and not abandoned, as to be likely, when used on or in connection with the goods of the applicant, to cause confusion, or to cause mistake, or to deceive . . . *Id.*

regarding search, we can assess the efficacy of various trademark search engines and study how machine learning methods can plausibly alter the landscape, potentially affecting trademark supply and quality.

Rather than focusing solely on the interaction between the consumer and the producer, our initial results suggest that AI can play a formidable role in addressing the cost of search regarding trademark selection, supply, and quality, warranting a greater focus on trademark producers and the registration ecosystem. While machine learning can minimize some preexisting search costs, our work shows that AI also carries the potential to introduce new search costs into the trademark ecosystem as well.

This work also carries implications for the economic literature regarding trademarks. Traditional approaches to trademarks, spearheaded by economic approaches, have focused almost exclusively on the demand-side role of search costs faced by the consumer. Yet we would argue that the economic literature on search costs, while valuable in considering consumer-based concerns, is incomplete in addressing various issues regarding trademark supply and quality. This conventional economic account fails to also consider the substantial search costs that are faced by not just consumers, but trademark applicants and firms as well in the process of trademark selection.

We argue, primarily, that in an age where AI will increasingly govern the process of trademark selection, this classic division between consumers and trademark owners needs updating, one which reflects that trademark applicants *also* function as consumers in the trademark selection ecosystem. In other words, rather than focusing on the relationships between trademark registrants and buyers or end users of products, we might also focus on how AI-powered search engines flip this dynamic and transform trademark applicants into consumers of trademarks as well. This insight, we suggest, has transformative potential for encouraging both innovation and efficiency in the process of trademark registration. In addition, it also suggests the need to study ways to deploy AI to better optimize search functions, thereby affecting trademark quality and the overall ecosystem as a result.

This Article has four parts. Part I outlines the basic contours of the traditional, demand-side approach in the economic literature focusing on consumer search costs in justifying trademark protection. Part II turns to introducing the role of AI in trademark search, explaining the legal and economic significance of a search cost theory that focuses on trademark supply, rather than demand. Part III turns to our empirical investigation, offering a comparison and contrast of various search engines to demonstrate how supply-side search considerations represent an important aspect of trademark theory. Finally, in Part IV we discuss the legal and economic

implications of our research, further exploring the potential role of AI in our legal system for trademarks.

I. SEARCH COSTS IN TRADEMARK LAW: A VIEW FROM THE CONSUMER

Back in 1961, George Stigler changed the field of consumer-related economics when he set forth a framework to understand the economic role of information in consumer decision-making.¹⁶ “One should hardly have to tell academicians that information is a valuable resource: knowledge is power,” he wrote.¹⁷ “And yet it occupies,” he wrote, “a slum dwelling in the town of economics.”¹⁸ Yet, if we consider the economic implications of the search for information in the market for goods, he predicted, we can better understand how it affects market price.¹⁹

Stigler’s insight—and the resulting body of literature that followed from it—has come to embody the “informative” view of advertising, one of the dominant approaches to an economic study of advertising.²⁰ Under this view, which originated out of the Chicago school in the 1960s, consumers often encounter search costs that deter them from learning about a product’s availability, price, and quality.²¹ Yet advertising, economists argue, can reduce the search costs for this information, improving the efficiency of the marketplace.²² As we show below, this general view has translated into a specific declaration of the economic and informative value of trademarks in this consumer-centric process of decision-making, a factor that lays the groundwork for a deeper examination of the centrality of search costs in the process of trademark selection.

16. See generally George Stigler, *The Economics of Information*, 69 J. POL. ECON. 213 (1961); see also Cathy Roheim Wessells, *The Economics of Information: Markets for Seafood Attributes*, 17 MARINE RES. ECON. 153, 154–55 (discussing Stigler).

17. Stigler, *supra* note 16, at 213.

18. *Id.*

19. *Id.*

20. See generally KYLE BAGWELL, *THE ECONOMIC ANALYSIS OF ADVERTISING* 6 (2005) (discussing the informative, persuasive, and complementary view of advertising).

21. *Id.* at 3.

22. See generally William M. Landes & Richard A. Posner, *Trademark Law: An Economic Perspective*, 30 J. L. & ECON. 265 (1987); Nicholas S. Economides, *The Economics of Trademarks*, 78 TRADEMARK REP. 523 (1988).

A. SEARCH, EXPERIENCE, AND CREDENCE ATTRIBUTES IN CONSUMER DECISION-MAKING

Traditional neoclassical economic theory implied that price signals convey all of the information necessary for consumers to make decisions.²³ However, today only a few markets reflect this phenomenon, because not only are most goods heterogeneous (offering a range of product attributes), but some of those attributes are observable, and others are not.²⁴ As a result, consumers make their decisions in a world of substantial information asymmetry. However, economists explain, advertising (and relatedly, trademarks) can reduce the costs of obtaining that information.²⁵ In turn, by offering protection to trademarks, the law thus reduces the search costs consumers face.

By reframing consumer decision-making to include a focus on the willingness to pay for information and the costs of obtaining it, Stigler opened up a world of greater inquiry on how producers communicate information to the public, and the implications of the cost of that information. Years later, in an influential set of papers, Philip Nelson refined Stigler's pathbreaking work by pointing out that there were even greater difficulties associated with ascertaining product quality than price, since information about quality is often impossible to discover before purchase.²⁶ This view of the consumer's asymmetric search for information has led to the classification of search and experience goods, a framework that underscores the function of trademarks in each category of the marketplace.²⁷ Others, including Ariel Katz, have since

23. See generally Jie "Jennifer" Zhang, Xiao Fang & Olivia R. Liu Sheng, *Online Consumer Search Depth: Theories and New Findings*, 23 J. MGMT. INFO. SYS., 72 (2006) ("Existing economic theory modeled consumers' search behavior as a compromise of the anticipated utility gain through price reduction and the additional search cost. Those models assumed that consumers are only searching for a single attribute (e.g., price).").

24. See generally *id.*

25. *Id.* at 82–83 (citing George A. Akerlof, *The Market for "Lemons": Quality Uncertainty and the Market Mechanism*, 84 Q. J. ECON. 488 (1970)); see also Landes & Posner, *supra* note 22, at 269 ("Rather than investigating the attributes of all goods to determine which one is brand X or is equivalent to X, the consumer may find it less costly to search by identifying the relevant trademark and purchasing the corresponding brand.").

26. See Wessells, *supra* note 16, at 155 (discussing Nelson); see generally Phillip Nelson, *Advertising as Information*, 82 J. POL. ECON. 729 (1974) (discussing that there are some qualities of a product which cannot be successfully conveyed by advertising).

27. Phillip Nelson articulated the distinction between search and experience goods; Darby and Karni added a third category, credence goods, to the mix. See Phillip Nelson, *Information and Consumer Behavior*, 78 J. POL. ECON. 311, 312 (1970); Michael R. Darby & Edi Karni, *Free Competition and the Optimal Amount of Fraud*, 16 J. L. & ECON. 67, 68–69 (1973).

pointed out that a more precise term might refer to these categories as “attributes,” instead of “goods.”²⁸

Each category nevertheless illustrates the importance of trademarks and advertising in ameliorating the information asymmetry faced by the consumer.²⁹ Search attributes are qualities that have characteristics which are observable to the consumer, and the brand or producer matters less because the product is readily identifiable (like, for example, table salt).³⁰ However, in the context of experience attributes, quality can only be determined after consumption of the good, like a newspaper or a law review article that needs to be read first for a consumer to determine its quality.³¹ Advertising and trademarks can improve the market for both search and experience attributes because they can provide consumers with pre-purchase information about both price and quality. This, in turn, has the effect of lowering consumers’ search costs in reaching decisions.³²

Later, economists added credence attributes as a third category. These involve goods like pharmaceuticals or automobile repair, where the quality cannot be determined until long after the good has been purchased and consumed.³³ Compared to search attributes and experience attributes, credence attributes are often infeasible to judge even right after purchase, and may take more time to ascertain their quality.³⁴ Thus, labeling and disclosure-related information can transform a credence attribute into a search attribute in order to empower a consumer to judge the quality of a good prior to making

28. Ariel Katz, *Beyond Search Costs: The Linguistic and Trust Functions of Trademarks*, 2010 BYU L. REV. 1555, 1561. We use the terms interchangeably although we note that Katz is correct that attributes is a more precise formulation.

29. *See id.* at 1560–61. Later, Nelson separated products into two different types: search goods and experience goods. *See* Nelson, *Consumer Behavior*, *supra* note 27 (exploring the ways by which a consumer acquires information about the quality of goods); *see also* Darby & Karni, *supra* note 27, at 68–72 (discussing the importance of credence attributes in assessing the value of the product); George Akerlof, *The Market for “Lemons”: Quality Uncertainty and the Market Mechanism*, Q.J. ECON. 488 (1970).

30. Katz, *supra* note 28, at 1560 (“For most consumers, all salt is equally salty, and as long as the consumer can reliably identify the white crystals as salt, the identity of the manufacturer or the exact brand chosen makes very little difference.”).

31. *Id.*; *see* Nelson, *Consumer Behavior*, *supra* note 27, at 312.

32. Wessells, *supra* note 16, at 155 (discussing Nelson).

33. Katz, *supra* note 28, at 1561. As Cathy Wessells pointed out, the markets surrounding credence goods are deeply imperfect. This is for two reasons: (1) because of the asymmetry of knowledge between the producer and the consumer and (2) because it is not practical or often even possible for consumers to assess the quality of the product beforehand (e.g., by performing laboratory tests, etc.). Wessells, *supra* note 16, at 155.

34. Darby & Karni, *supra* note 27, at 69.

a purchase.³⁵ Consumers of such goods may perhaps also be aided by a certification of a good by an external source.³⁶ “For credence goods,” Cathy Wessells writes, “one may rely on producer claims, but generally consumers place more trust in an independent third party to provide truthful information on quality,” suggesting a role for independent third-party private certification (i.e., certification trademarks) or government regulation.³⁷

Taken together, these categories of goods appeared in a substantial amount of economic and legal literature on the foundational role played by advertising—and trademarks—in addressing consumer decision-making. Trademarks, like other forms of advertising, provide important information to both consumers and other producers about their source.³⁸ In search and experience goods, advertising minimizes the information asymmetry faced by the consumer, enabling her to process information about the good and to decide whether or not to purchase. As Nicholas Economides explains, “Where experience goods have unobservable differences in quality and/or variety, trademarks enable consumers to choose the product with the desired combination of features and encourage firms to maintain consistent quality and variety standards and to compete over a wide quality and variety spectrum.”³⁹ In other words, trademarks convey valuable information for all three categories of attributes, thus justifying their legal protection.

B. TRADEMARK LAW AND CONSUMER SEARCH COSTS

The above analysis describes the role played by trademarks in identifying each of the three core categories of attributes, thereby reducing consumer search costs. Even the Supreme Court has endorsed the search cost justification for trademark protection.⁴⁰ In *Qualitex*, the Court noted,

35. Wessells, *supra* note 16, at 155 (citing Caswell). See Julie A. Caswell, *Valuing the Benefits and Costs of Improved Food Safety and Nutrition*, 42 AUSTL. J. AGRIC. & RES. ECON. 409 (1998).

36. Darby & Karni, *supra* note 27, at 69–70 (outlining credence goods, by taking the example of repair services, which basically requires a consumer to purchase both information (about the diagnosis of, say, a malfunctioning machine) and repair (actual performance of the repair)). If there were no additional costs involved in separating the two then the authors suggest that the consumer would do so in order to avoid the possibility of fraud. But since it is often cheaper to provide information and service jointly, then the consumer will purchase them both from the same source.

37. Wessells, *supra* note 16, at 155.

38. Stacey L. Dogan & Mark A. Lemley, *Trademarks and Consumer Search Costs on the Internet*, 41 HOUS. L. REV. 777, 777–78 (2004).

39. Economides, *supra* note 22, at 525.

40. See Mark P. McKenna, *A Consumer Decision-Making Theory of Trademark Law*, 98 VA. L. REV. 67, 75–76 (2012) (“The overwhelming majority of scholars use search costs language to describe trademark law’s purposes, and the Supreme Court has explicitly endorsed the theory as trademark law’s core theoretical justification.” (internal citation and quotations omitted));

“[T]rademark law, by preventing others from copying a source-identifying mark, ‘reduce[s] the customer’s costs of shopping and making purchasing decisions,’ for it quickly and easily assures a potential customer that this item—the item with this mark—is made by the same producer as other similarly marked items that he or she liked (or disliked) in the past.”⁴¹

Similarly, William Landes and Richard Posner frame trademarks primarily as an informational mechanism to provide consumers with information about the seller’s identity, the quality of the product, etc., and thereby reduce the consumer’s search costs for comparable goods.⁴² The search cost approach has had multiple implications for trademark law; among them are reinforcing the centrality of the consumer and also indirectly empowering strong marks over weaker ones.⁴³ As Barton Beebe has pointed out, the more distinctive the mark, the less costly it is for the consumer to locate in the marketplace; thus, stronger marks better facilitate the search process for consumers than weaker marks.⁴⁴ Trademarks also help guarantee market quality, ameliorating the market failure George Akerlof identified in his famous piece.⁴⁵ Not only do they reduce search costs by condensing complex information into an identifiable symbol, but they also “allow buyers to trust and rely upon the signals conveyed by

see also WILLIAM M. LANDES & RICHARD A. POSNER, *THE ECONOMIC STRUCTURE OF INTELLECTUAL PROPERTY LAW* 166–209 (2003); John F. Coverdale, *Trademarks and Generic Words: An Effect-on-Competition Test*, 51 U. CHI. L. REV. 868, 869–70, 878 (1984); Stacey L. Dogan & Mark A. Lemley, *A Search-Costs Theory of Limiting Doctrines in Trademark Law*, 97 TRADEMARK REP. 1223, 1223 (2007); Stacey L. Dogan & Mark A. Lemley, *Grounding Trademark Law Through Trademark Use*, 92 IOWA L. REV. 1669, 1689–90, 1697 (2007); Economides, *supra* note 22, at 525–27; Michael Grynberg, *The Road Not Taken: Initial Interest Confusion, Consumer Search Costs, and the Challenge of the Internet*, 28 SEATTLE U.L. REV. 97, 97–99 (2004); William M. Landes & Richard A. Posner, *The Economics of Trademark Law*, 78 TRADEMARK REP. 267, 272 (1988); Mark A. Lemley, *The Modern Lanham Act and the Death of Common Sense*, 108 YALE L.J. 1687, 1695–96 (1999); Clarisa Long, *Dilution*, 106 COLUM. L. REV. 1029, 1033–34, 1056 (2006); Glynn S. Lunney, Jr., *Trademark Monopolies*, 48 EMORY L.J. 367, 432 (1999); I.P.L. Png & David Reitman, *Why Are Some Products Branded and Others Not?*, 38 J.L. & ECON. 207, 208–11 (1995).

41. *Qualitex Co. v. Jacobson Prods. Co., Inc.*, 514 U.S. 159, 163–64 (1995) (internal citations omitted).

42. See Landes & Posner, *Trademark Law: An Economic Perspective*, *supra* note 22, at 269–70.

43. For an excellent account of the multiple roles of search in trademark law, see Barton Beebe, *Search and Persuasion in Trademark Law*, 103 MICH. L. REV. 2020, 2042 (2005).

44. *Id.* at 2042–43.

45. See Akerlof, *supra* note 29 (arguing that in situations where sellers and buyers have asymmetric information about the quality of a good (i.e., with a used car), adverse selection will occur where high-quality sellers leave the market as consumer willingness-to-pay falls). To avoid this type of market failure, building credible signals of product quality is crucial, and advertising can help achieve this goal.

sellers as guarantees for quality, thus helping to prevent the lemonization of markets for goods with experience and credence attributes.”⁴⁶

Firms that produce experience or credence goods are therefore incentivized to keep a consistent level of quality associated with their goods in order to ensure repeat purchasers; trademarks reduce search costs in both of these arenas, enabling the consumer to trust that the purchase they are making will be consistent with their prior experience.⁴⁷ But, as Mark Lemley and Stacey Dogan explain, there is a crucial catch: this only works if consumers can readily trust the information that trademarks provide, thereby paving the way for the role of law.⁴⁸ “By protecting established trademarks against confusing imitations,” they write, “the law ensures a reliable vocabulary . . . Both sellers and buyers benefit from the ability to trust this vocabulary to mean what it says it means.”⁴⁹ Because trademarks economize on information, it is thought that making it less costly to obtain will better inform consumers and thereby improve the competitiveness of the market.⁵⁰

Despite the potentially rich layers of focus on trademark owners and applicants for discussion, no other theory has managed to displace the primary importance of the search-cost rationale and its consumer-centric focus. Mark McKenna has valuably pointed out that trademark law itself predated the search cost theory by several hundred years, suggesting that a historical account might be a better, more comprehensive theory to address its development.⁵¹ Other scholars have written about how trademark protection performs a “signaling” function within advertising; others have focused on how brands facilitate corporate growth into new territories; and still others focus on how trademarks are viewed as a kind of property right.⁵² Yet, despite

46. Katz, *supra* note 28, at 1563.

47. Katz, *supra* note 28, at 1561. While these classes of goods are incredibly helpful in distilling the marketplace, Ariel Katz reminds us that in more contemporary parlance, it is more correct to refer to attributes instead of goods.

For example, the fact that a can of tuna looks like a can of tuna is a search attribute. The fact that the content tastes like tuna is an experience attribute.

Whether the content is indeed tuna and not a good imitation, or whether it is safe for consumption, are credence attributes.

Id. at 1561.

48. Stacey L. Dogan & Mark A. Lemley, *Trademarks and Consumer Search Costs on the Internet*, 41 HOUS. L. REV. 777, 786–87 (2004).

49. *Id.* at 787.

50. *Id.*

51. McKenna, *supra* note 40, at 67.

52. Dogan & Lemley, *Trademarks and Consumer Search Costs on the Internet*, *supra* note 48, at 799; see also Ralph S. Brown Jr., *Advertising and the Public Interest: Legal Protection of Trade Symbols*, 57 YALE L.J. 1165, 1184 (1948); Lemley, *The Modern Lanham Act and the Death of Common Sense*,

the promise of these alternative approaches, search cost theory still plays a seminal role in trademark law, often ensuring the consumer's centrality to trademark law, at times even at the expense of a trademark owner.

Multiple doctrines of trademark law—distinctiveness, genericness, dilution, comparative advertising, and even the theory of trademark use—implicitly follow the search cost approach in crafting legal entitlements.⁵³ For example, the goal of limiting search costs has been implicitly extended to explain the genericness doctrine, in order to avoid the risk that “[c]onsumers will be misled if what they believe is a generic term is in fact a product sold by only one company.”⁵⁴ The search cost rationale has also been extended to justify Congress's foray into enacting federal anti-dilution protections, under the reasoning that uses that blur or tarnish famous marks increase the search costs faced by the consumer by either weakening the meaning of the mark in the eyes of the consumer or creating a negative impression of or association with the mark.⁵⁵ In sum, trademarks have served as a vehicle to optimize consumer access to information through reducing search costs, and much of trademark law has integrated this goal throughout various doctrines.

II. SEARCH COSTS IN TRADEMARK REGISTRATION: A VIEW FROM A TRADEMARK APPLICANT

As we discussed above, the conventional legal accounts of search costs focus largely on improving the information shared with the consumer. But this view can often be too narrow. Very little attention is paid to the process of optimizing the information markets that develop around the process of trademark search and registration, even though these variables can have a dramatic effect on trademark supply and enforcement.⁵⁶ However justifiable

supra note 40, at 1714; Kenneth L. Port, *Trademark Monopolies in the Blue Nowhere*, 28 WM. MITCHELL L. REV. 1091 (2002); Lunney, Jr., *Trademark Monopolies*, *supra* note 40; Frank I. Schechter, *Fog and Fiction in Trade-Mark Protection*, 36 COLUM. L. REV. 60, 65 (1936).

53. See Dogan & Lemley, *supra* note 48, at 786–99.

54. At the same time, however, Lemley and Dogan point out that the genericness doctrine can actually increase search costs if an ultra-famous mark like “aspirin” or “thermos” has now become generic, since consumers who might associate the mark with a particular source may now be confused if the term is used to refer to a class of goods instead. See *id.* at 793.

55. *Id.* at 789–90; see also Rebecca Tushnet, *Gone in Sixty Milliseconds: Trademark Law and Cognitive Science*, 86 TEX. L. REV. 507 (2008) (noting the argument, aided by cognitive science, that negative trademarks (either ones that weaken or tarnish a mark) can create informational harms that reduce consumers' capacity to shop around in a rational manner).

56. Of course, see the seminal paper by Beebe and Fromer, which valuably focused on the issue of trademark supply. See Barton Beebe & Jeanne C. Fromer, *Are We Running Out of*

the search cost approach may be, it can affect the trademark supply if it adds too much strength to established marks at the cost of others. Too much empowerment of trademark holders can enable them to exert overbroad control over uses that may not even be legitimate trademark uses, or to stifle competitors who are simply describing their own products.⁵⁷ As Lemley and Dogan point out, stronger trademark entitlements can also have the effect of narrowing the scope of available words for others to use.⁵⁸

Moreover, despite all of the analysis surrounding the consumer, there is very little recognition of the fact that trademark registrants are also consumers as well in the marketplace of trademark search and registration. Even aside from the law's role in registration, the selection of a trademark is a crucial moment for a firm because it symbolizes much more than the source of the product itself. Since the goal of modern marketing and branding is to essentially create desire among consumers by making irrelevant attributes seem relevant and valuable,⁵⁹ the selection of an appropriate trademark is an emotionally-driven choice as well as an economic one.⁶⁰ Brands confer market power. As one author writes, "when trademarks protect brands with significant image value, the brand in and of itself becomes a product characteristic that consumers care about but competitors cannot copy."⁶¹

Thus, the same price and non-price variables that might influence a consumer's purchasing decision might also influence a trademark registrant's decision to select a mark. Even information about the demographics of the typical and non-typical trademark registrants and their trademark search processes or sophistication with online search would be enormously helpful in future research.⁶² AI-driven tools could play a crucial role in this process at all levels ranging from trademark selection, to application, and to registration.

Moreover, in a world characterized by more trademarks than ever, it becomes necessary to explore the costs incurred by firms themselves in the process of searching for available trademarks. Trademark applicants will

Trademarks? An Empirical Study of Trademark Depletion and Congestion, 131 HARV. L. REV. 945, 947 (2018).

57. Dogan & Lemley, *supra* note 48, at 788.

58. *Id.*

59. See McKenna, *supra* note 40, at 115 (citing Gregory S. Carpenter et al., *Meaningful Brands from Meaningless Differentiation: The Dependence on Irrelevant Attributes*, 31 J. MKTG. RES. 339, 339 (1994)).

60. WORLD INTELLECTUAL PROP. ORG., *supra* note 5, at 86. See generally Sonia Katyal, *Stealth Marketing and Antibranding: The Love that Dare not Speak its Name*, 58 BUFF. L. REV. 58 (2010) (discussing the lure of branding); Sonia Katyal, *Trademark Cosmopolitanism*, 47 UC DAVIS L. REV. 875 (2013) (discussing the emergence of brands as global figures of speech).

61. WORLD INTELL. PROP. ORG., *supra* note 5, at 86.

62. See Zhang et al., *supra* note 23, at 91 (noting the role of similar attributes for a typical study of consumer search behavior).

expend tremendous effort and incur costs in order to find their optimal trademark for both economic and non-economic reasons. These kinds of search costs seem to be underexplored in the relevant trademark literature, but they are important. Because of the economic benefits of maintaining trustworthy trademarks, the USPTO will reject trademark applications that risk trademark infringement or dilution. To avoid this risk, a firm will ideally want to avoid the costs associated with filing a doomed application, and instead preemptively search for existing marks and calculate the probability of infringement or dilution based on those search results. For this reason, AI and machine learning can play a significant role in improving trademark quality and registrability, reducing the search costs faced by trademark applicants.⁶³

Below, we outline the theoretical basis for studying how private AI-powered search tools have emerged to play an important role in supplementing government determinations and reducing search costs faced by the trademark applicant. We then turn to the specifics of discussing how AI is used by government agencies in administering IP and by private entities in the process of search, registration, and brand management.

A. SUPPLEMENTING TRADEMARK SEARCH IN THE PRIVATE SECTOR

In Part I, we discussed the traditional economic underpinnings of trademarks from the consumer's point of view. Specifically, we discussed the need for the USPTO to avoid granting marks that would result in informational harms to consumers. An erroneously granted trademark creates harms to consumers by confusing them and eroding their ability to discern meaningful information about a good or service. In turn, this situation would harm the original holder of a trademark that relies on the guarantee of quality that their mark provides in order to sell their products to consumers. But even before the PTO makes its determination, machine learning can also help to optimize the search process from an applicant's perspective, thus providing a role that essentially supplements the PTO's eventual determination by lowering the search costs associated with trademark selection.

While this paper is concerned with the deployment of machine learning in trademark search and registration, it is important to note that a few scholars

63. WORLD INTELL. PROP. ORG., *supra* note 5, at 107. Outside of the trademark law community, there is a robust conversation ongoing about the future uses of AI for both litigation and transaction-related tasks. See John Markoff, *Armies of Expensive Lawyers, Replaced by Cheaper Software*, N.Y. TIMES (Mar. 4, 2011), <https://www.nytimes.com/2011/03/05/science/05legal.html>; see also Timothy J. Carroll & Manny Caixeiro, *Pros and Pitfalls of Artificial Intelligence in IP and the Broader Legal Profession*, LANDSLIDE (Jan. 2019), <https://www.dentons.com/en/-/media/fa72a6d5cb304c1194e015eb26123e27.ashx>.

have analyzed its use in patent applications.⁶⁴ In a thoughtful piece about machine learning at the USPTO, Arti Rai discusses the use and implications of its impact in the area of prior art search, noting that it holds significant promise in maximizing efficiency in a world of overburdened patent office administration.⁶⁵ While Rai focuses much of her analysis on USPTO reliance on machine learning, her work valuably opens up a larger discussion about the relationship between AI-driven private search engines and the USPTO's own tools.

Both Rai and Tabrez Ebrahim⁶⁶ have noted that AI tools enable patent applicants to design their applications in a way that maximizes their information advantages.⁶⁷ Patent applicants have private information about the quality and originality of their patents, and patent examiners must work to uncover this information and make decisions about patentability.⁶⁸ Ebrahim valuably explores this idea of information asymmetries between the patent office and the private sector at length. In a model, described as the Spence Model of Information Exchange, he describes a back-and-forth game where the patent applicant and patent office engage in countering signals about the patent's quality.⁶⁹ The applicant is always the first mover and will try to maximize the scope of the patent application, and the patent examiner tries to discern whether this scope is reasonable and may try to pare it back.⁷⁰ The examiner and patent applicant (or the patent prosecutor) will go back and forth until they settle on an equilibrium.⁷¹ Ebrahim argues that success in this game rests on each party's ability to discover relevant information.⁷²

Critically for our study, he also describes how privately supplied AI tools can exacerbate information asymmetries between the patent applicant and the

64. See, e.g., David Engstrom, Daniel E. Ho, Catherine M. Sharkey & Mariano-Florentino Cuéllar, *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies* 46–52 (2020), available at <https://www-cdn.law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf>.

65. Rai, *supra* note 3, at 2619–21; see generally Michael D. Frakes & Melissa F. Wasserman, *Irrational Ignorance at the Patent Office*, 72 VAND. L. REV. 975 (2019) (concluding that each patent examiner needs more time to assess a patent application to improve patent quality); U.S. General Accountability Office, *Intellectual Property: Patent Office Should Strengthen Search Capabilities and Better Monitor Examiners' Work*, GAO-16-479 (July 20, 2016), <https://www.gao.gov/products/GAO-16-479> (recommending steps to improve the prior art search quality).

66. Ebrahim, *supra* note 3, 104.

67. *Id.* at 1196–1201.

68. *Id.* at 1211–12.

69. *Id.* at 1191.

70. *Id.*

71. *Id.*

72. *Id.* at 1221–23.

patent office because the patent office does not have the tools to discern between high- and low-quality signals.⁷³ Thus, the patent office will be in a position where it cannot adequately sift through a market for lemons, thus creating a supply-side issue where the generators of information can more successfully play the information game.⁷⁴ More broadly, AI could also displace the need for lawyers, as he explains that:

[a]rtificial-intelligence technology could displace or reduce the need for attorneys in law firms or in-house legal departments and, in doing so, lessen the job opportunities for law students. The impact of decreasing the role of legal-service professionals with AI technology affects the relationship between clients and lawyers and, as a result, also affects the relationship of the interaction between inventors and the USPTO.⁷⁵

We might imagine that similar forces are at play with trademarks. Although trademark approvals, particularly simple word marks, are likely not as complex as patent examinations, there is evidence that AI is transforming this area of IP law as well. The impact of AI on trademark search may be greatest for word marks or composite marks with literal elements, since more data might be available, allowing for greater ease of identifying similarities and differences.⁷⁶

In essence, however, the core search cost problem that Ebrahim and Rai articulate from the perspective of patent applicants and examiners is the same problem that we are exploring from the perspective of trademark applicants. The rise of the private sector in search can have dramatic effects on trademark quality and supply, just like in the patent context. Primarily, the “likelihood of confusion” standard in trademarks is similar to the non-obviousness standard in patents because of the human subjectivity involved in both processes. Each requires an examiner determining whether to grant an application based on their best evaluation of the application, with an eye toward minimizing errors that could result in informational harms to consumers.

Here, we might also note the risk that private vendors’ search tools might be more sophisticated than those of the government.⁷⁷ Indeed, the emergence

73. *Id.* at 1220.

74. *See id.* at 1236.

75. *Id.* at 1231–32.

76. *Letter from American Bar Association-Intellectual Property Law Section to Secretary of Commerce for Intellectual Property & Director of the United States Patent and Trademark Office*, USPTO (Jan. 9, 2020), https://www.uspto.gov/sites/default/files/documents/ABA-IPL_RFC-84-FR-58141.pdf, at 12 [hereinafter ABA Letter].

77. She also discusses the risks in relying on private vendors from an explainability/due process perspective, observing that there is at least an appreciable risk that using private search

of a private market for trademark search indicates that there may be a market failure regarding trademark registration. Although the USPTO operates its own free search service, there are several private sector alternatives.⁷⁸ These private services variously advertise their added value as being powered by AI, machine learning, statistical models, or other sophisticated techniques.⁷⁹ Insofar as trademark applicants rely on these private services instead of the USPTO, it suggests that these services provide real value that the government service does not.⁸⁰

Moreover, since the USPTO is not an enforcement agency, and IP rights owners are responsible for protecting their marks, the government may not have the right incentives to have the best AI tools available, and can instead externalize these costs to trademark registrants. This externalization thus creates a market for the sorts of private AI tools in our study, which function to supplement the government's inadequate TESS system. Assuming that the USPTO relies on its own TESS search engine, and that TESS does not work as well as these AI-powered private sector alternatives, the emergence of private search engines suggests that the government's inadequacy may be potentially (indirectly) imposing costs on trademark holders and consumers.

An increase in AI-powered search could plausibly reduce the overall number of applications filed because it would forecast which marks were likely to face a Section 2(d) refusal.⁸¹ Consider: both examiners and applicants want to avoid the monetary and time costs associated with bad applications. A trademark can cost about \$250 per class it is registered for,⁸² and it takes a substantial amount of time.⁸³ While the cost of the mark may be trivial for larger companies and brands, the time involved and attorney's fees can be

vendors might result in assertions of trade secrecy and more opacity. Rai, *supra* note 3, at 2640–41.

78. For example, see Corsearch, Markify, Trademarkia, and TrademarkNow. We detail these in a below section.

79. See Nick Potts, *Reviews of the 3 Best Trademark Clearance Search Tools for Trademark Attorneys*, TRADEMARKNOW (Oct. 20, 2016), <https://www.trademarknow.com/blog/reviews-of-the-3-best-trademark-search-tools-for-trademark-attorneys>.

80. Part of this extra value-added may come from the fact that the AI technologies underlying trademark search are also used for brand protection. We discuss this further in Part III.

81. See ABA Letter, *supra* note 76, at 11–12.

82. U.S. Patent & Trademark Office, *Trademark Fee Information*, <https://www.uspto.gov/trademark/trademark-fee-information> (last visited on Jan. 22, 2021).

83. See U.S. Patent & Trademark Office, *Section 1(b) Timeline: Application Based on Intent to Use your Trademark in Commerce*, <https://www.uspto.gov/trademark/trademark-timelines/section-1b-timeline-application-based-intent-use> (last visited on Jan. 22, 2021).

substantial.⁸⁴ The USPTO provides a useful chart, included as Figure 1, for a 1(b) trademark application—essentially when an applicant files a mark with intent to use it later.⁸⁵ At a minimum, from the time an application is filed to when it is approved is about seven months. However, if the USPTO does not immediately approve the mark, it adds at least three months to the process, and as much as an additional eight months if there are multiple rounds of correspondence between the applicant and the USPTO.⁸⁶ That additional time could represent lost revenue and other harms stemming from lack of IP protection.

From the USPTO's point of view, AI might provide assistance in achieving greater consistency among Examining Attorneys by helping them reach faster decisions, reducing their workload, and enabling them to identify any inconsistencies in outcomes.⁸⁷ It might also aid the detection of fraudulent filings and practices as well, through its evaluation of metadata and closer image comparisons.⁸⁸ If AI can be used by applicants to ensure that they do not erroneously file an application that is destined to undergo additional rounds of screening or a final rejection from the USPTO, they can save the time and energy needed to go through the appeals process.

84. The examination process involves three steps: first, the mark is classified into a series of design codes; second, examiners search through existing marks, pending applications, and abandoned marks for similarity; and third, issue a determination regarding whether the mark is eligible for registration. *See* Engstrom et al., *supra* note 64, at 47.

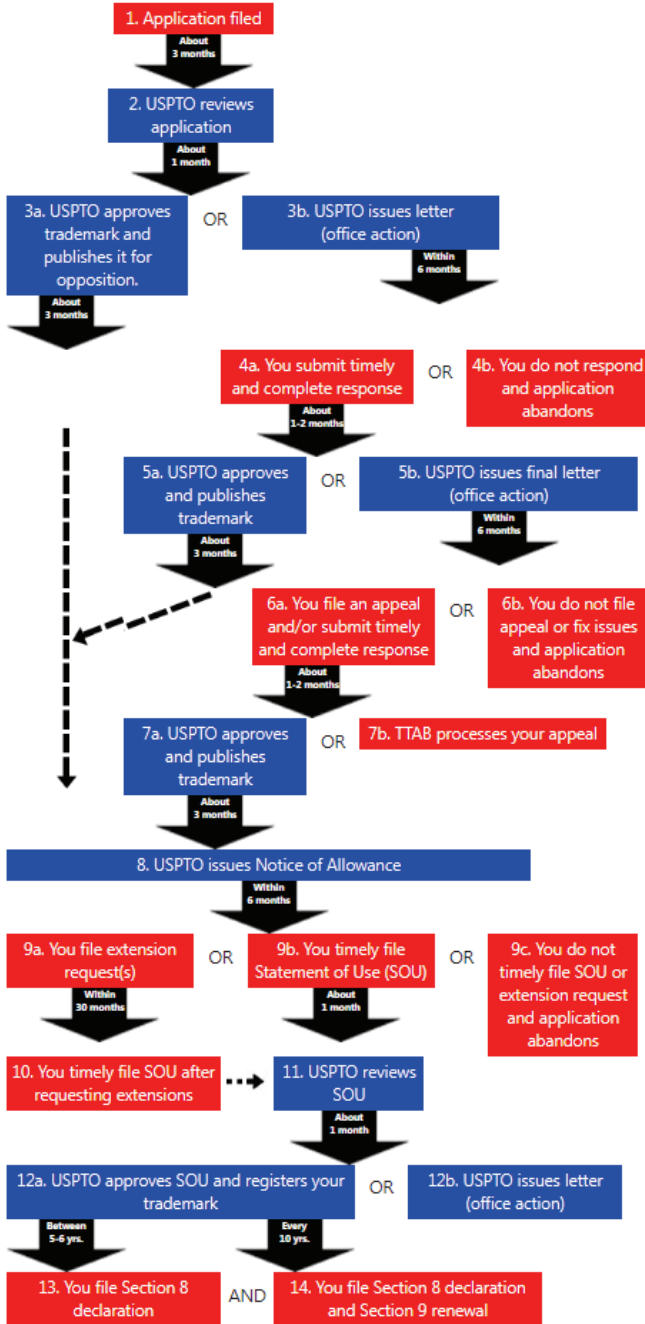
85. *Id.*

86. *See Section 1(b) Timeline*, USPTO, <https://www.uspto.gov/trademark/trademark-timelines/section-1b-timeline-application-based-intent-use> (last visited on Jan. 22, 2021) (Figure 1 below) (showing the timeline for 1(b) applications).

87. *See* ABA Letter, *supra* note 76, at 12.

88. *See id.*

Figure 1: Timeline of Section 1(b) Applications



B. ARTIFICIAL INTELLIGENCE AND THE ADMINISTRATION OF
INTELLECTUAL PROPERTY

The prior Section provided a general theoretical basis for the emergence of a private search market for trademarks. A key part of this inquiry involves differentiating between (1) government use of AI-powered techniques to assist with their determinations; (2) AI-powered tools made available by the government to assist private parties in preparing applications for patent, trademark, or copyright protection; and (3) a comparably broader set of AI-driven private tools developed for private parties (rather than the government) in order to supplement state-offered techniques. Below, we focus on the first two categories, in discussing the role of AI in government-led administration of IP and explaining how this paves the way for a private AI-powered market to optimize trademark search and registration. We turn to the third category in our next section.

Last year, WIPO released the first comprehensive global survey of how AI and machine learning can be employed to assist with the governance of IP.⁸⁹ Out of WIPO's survey of thirty-five different IP Offices, the report noted that seventeen offices use AI technology in at least one aspect of their work, but that most of these uses appear to be in their infancy.⁹⁰ On the patent side, the report points out that AI can be used to automatically analyze the content of patent applications and case files including sorting and allocating them for particular staff, as well as for applying particular classifications.⁹¹ It can also be used for the purposes of searching for prior art and to improve detection of links between citations and applications,⁹² and even to assist in the processing of applications.⁹³ One office in Singapore estimated saving five thousand hours of an examiner's man-hours by relying on AI techniques.⁹⁴

The USPTO is using machine learning in its determinations of patentability and histories of patent prosecution.⁹⁵ The USPTO, for example, developed a tool named Sigma, which can search an entire patent document

89. See *Meeting of Intellectual Property Offices (IPOS) on ICT Strategies and Artificial Intelligence (AI) for IP*, WORLD INTELL. PROP. ORG. (May 23–25, 2018), https://www.wipo.int/edocs/mdocs/mdocs/en/wipo_ip_itai_ge_18/wipo_ip_itai_ge_18_1.pdf.

90. See *id.*

91. *Id.* at 4 (noting developments in Germany, Brazil, and Singapore along these lines).

92. *Id.* at 5–6 (describing patent search systems).

93. *Id.* at 6 (describing Tequmine in Finland for patent classification and prior art search).

94. *Id.* at 12.

95. Isi Caulder & Paul Blizzard, *Canada: Artificial Examiner: The Expanding Use of AI and ML Software at Intellectual Property Offices (IPOS)*, BERESKIN & PARR LLP (July 26, 2018), <https://www.bereskinparr.com/doc/artificial-examiner-the-expanding-use-of-ai-and-ml-software-at-intellectual-property-offices-ipos>.

and compare applications with registered patents and pre-grant publications.⁹⁶ Rai noted that Sigma enables examiners to attach a particular weight to the most relevant part of the patent application, and then retrieve related documents including related prior art.⁹⁷ Another cluster of AI applications are also used to manage IP files' prosecution and formality checks, particularly regarding data support, proofreading, and conversion of files to enhance machine-readability, and also for the purposes of translation and data analysis.⁹⁸ According to one study led by David Engstrom, the USPTO is also considering ways to build an AI-driven search platform that would use content-based engines to suggest prior art for an applicant; other plans involve using neural word embeddings to expand prior art searches.⁹⁹

On the trademark side, as opposed to patent, the main focus so far has been on the process of search, which historically has been mostly manual.¹⁰⁰ In the context of trademark search, new tools would provide a valuable service by helping potential registrants identify potential conflicts before ever submitting a trademark application, providing the statistical tools necessary to distinguish signal and noise. Colleen Chien, echoing Ebrahim and Rai, has pointed out that the USPTO itself has a difficult time assessing patent quality and frequently grants patents that it probably ought not to.¹⁰¹ The same might also be said of trademarks, which compels the employment of AI-driven tools in searching for similar marks.

In the context of trademarks, AI has mainly been deployed as an enhancement tool to assess trademark similarity; however, as Dev Gangjee has noted in an excellent study, AI can play a broader, potentially game-changing role.¹⁰² In the government context, it is unlikely that AI will replace human judgment regarding the more complex and subjective tests in trademark law; however, AI still carries the power to streamline administrative tasks relating to registration, opposition, and other procedures, and is likely to only grow in importance.¹⁰³

Traditional search systems employ text-based retrieval technology; today, the technology has improved in order to incorporate phonetic analogies, synonyms, and related permutations of letters in order to compare slightly

96. See WORLD INTELL. PROP. ORG., *supra* note 89, at 3.

97. Rai, *supra* note 3, at 2634.

98. WORLD INTELL. PROP. ORG., *supra* note 89, at 2, 9–10 (noting developments in Singapore, China, Japan, Morocco, Serbia, and Canada).

99. Engstrom et al., *supra* note 64, at 48.

100. *Id.*

101. See Colleen V. Chien, *Comparative Patent Quality*, 50 ARIZ. ST. L.J. 71, 72–74 (2018).

102. Gangjee, *supra* note 6, at 2.

103. Moerland & Freitas, *supra* note 6, at 27.

modified marks as well.¹⁰⁴ Even more, AI has driven significant advances in three additional dimensions of search and comparison: (1) text and conceptual similarity (i.e., assessing text, as well as shared or oppositional meanings of a trademark);¹⁰⁵ (2) visual/image similarity (i.e., assessing image elements of a trademark logo or figurative mark, including content-based image retrieval);¹⁰⁶ and a combination of words and images in order to integrate both in a similarity assessment.¹⁰⁷ Still other approaches rely on a constellation of comparisons—such as automated similarity assessments of image/pixel, text, and content, coupled with a manual comparison—in order to provide a more comprehensive comparison.¹⁰⁸

There is growing evidence of government use of these AI-driven tools as well. Reports indicate that current government uses of AI in the context of trademarks involve image recognition, classification of goods and services, and identifying descriptive terms.¹⁰⁹ According to David Engstrom, the USPTO is also prototyping a deep learning model that uses an unsupervised approach to

104. Gangjee, *supra* note 6, at 6 (citing C.J. Fall & C. Giraud-Carrier, *Searching Trademark Databases for Verbal Similarities*, 27(2) WORLD PATENT INFO. 135 (2005)).

105. Gangjee, *supra* note 6, at 6–7 (advances in search technology based on semantic or conceptual similarity focus more on “lexical relations,” integrating assessments of synonyms, antonyms, or comparable words in another language) (citing F. Mohd Anuara, R. Setchia & Y-K Lai, *A Conceptual Model of Trademark Retrieval based on Conceptual Similarity*, 22 PROCEDIA COMPUT. SCI. 450, 451 (2013)).

106. Gangjee, *supra* note 6, at 7 (noting that WIPO and the European Intellectual Property Office offer users the ability to upload image-based file formats). Currently, WIPO relies upon a system, the International Classification of the Figurative Elements on Marks, also called the Vienna Classification system. Trademark examiners, in general, manually index and code elements of figurative marks, often in reference to the Vienna Classification system, and then match the Vienna codes of a new application with those already registered. Since not all trademark registries use the system, and it involves some subjectivity, there is the risk of gaps in its application. *See id.* at 7–8 (citing WIPO, *Future Development of the Vienna Classification: Questionnaire Results* (April 3, 2019)). According to Gangjee, AI-assisted processes of content-based image retrieval have been “welcomed,” due to the added value of accuracy in comparison. *Id.*

107. Gangjee, *supra* note 6, at 6–9. As he writes, “[t]he goal is to mimic the assessment of a human examiner who must synthesize visual, aural, and conceptual similarity to arrive at an overall conclusion on whether the marks conflict.” *Id.*

108. *Id.* at 10 (citing Mosseri I., Rusanovsky M. & Oren G., *TradeMarker – Artificial Intelligence Based Trademarks Similarity Search Engine*, in COMMUNICATIONS IN COMPUTER AND INFORMATION SCIENCE (vol. 1034, 2019), https://doi.org/10.1007/978-3-030-23525-3_13); Moerland & Freitas, *supra* note 6, at 2 (noting that only a few trade mark offices apply AI tools).

109. Moerland & Freitas, *supra* note 6, at 15; *see also* Engstrom et al., *supra* note 64, at 49 (describing the use of a deep learning image classifier and other prototypes).

generate visually similar images from a database.¹¹⁰ The International Trademark Association has reported that at least five governments have developed trademark image search engines that incorporate AI.¹¹¹ The USPTO, for example, has developed a manually coded system of figurative images in order to train its deep learning systems to generate design codes for new trademark image applicants.¹¹² Other governments rely more extensively on private image search tools for their government registries.¹¹³ For example, IP Australia and the E.U. Intellectual Property Office uses TrademarkVision's Image Recognition (now a part of Clarivate Analytics) to search existing trademark images, employing a technology similar to facial recognition technology, but applied to marks instead.¹¹⁴ Chile, China, and Japan also rely on private tools.¹¹⁵ Some offices, such as that of Australia, even offer the public a range of AI-driven tools to assist unregistered applicants.¹¹⁶ And WIPO's Global Brand Database recently released a free AI-driven image search tool for the public.¹¹⁷

The wide range of emerging tools may lead some to suggest that AI might even have the effect of shrinking the potential role of the trademark lawyer. Since automated technologies can play a wider role in brand clearance and brand protection, it would enable service providers to work directly with trademark owners themselves. Echoing this view, others have observed that AI's added efficiency has the potential to replace paralegals or junior lawyers, perhaps when it comes to search and registration.¹¹⁸ However, more complex situations still call for human intervention. One WIPO survey respondent from Norway was careful to note that in comparing AI and non-AI results, while the most "similar" trademarks often had the same results, there were

110. Engstrom et al., *supra* note 64, at 49–50 (also describing future ways to deploy AI in image/text classification).

111. See *INTA Comments in Response to Request for Comments on Intellectual Property Protection for Artificial Intelligence Innovation*, USPTO 1, [https://www.uspto.gov/sites/default/files/documents/International%20Trademark%20Association%20\(IN_RFC-84-FR-58141\).pdf](https://www.uspto.gov/sites/default/files/documents/International%20Trademark%20Association%20(IN_RFC-84-FR-58141).pdf) (noting that out of 9 respondents to its survey, five IP offices are using AI-driven tools in trademark image search systems).

112. Gangjee, *supra* note 6, at 9 (citing U.S. Patent Trademark Office, *Emerging Technologies in USPTO Business Solutions* (May 25, 2018), https://www.wipo.int/edocs/mdocs/globalinfra/en/wipo_ip_itai_ge_18/wipo_ip_itai_ge_18_p5.pdf).

113. Gangjee, *supra* note 6, at 9.

114. *Id.*

115. WORLD INTELL. PROP. ORG., *supra* note 89, at 7–8 (describing developments).

116. *Id.*

117. See *Global Brand Database*, WORLD INTELL. PROP. ORG., <https://www3.wipo.int/branddb/en/> (last visited July 8, 2020).

118. See *How AI Impacts Trademarks*, TRADEMARK TIMES 1 (2018), <https://www.managingip.com/pdfsmip/01-TrademarkTimes18Seattle.pdf>.

very large differences found between AI and non-AI results in addressing lower degrees of similarity.¹¹⁹ This suggests that a mix of human and non-human intervention and greater amounts of data would improve the outcome.¹²⁰

Beyond image search, offices reported relying on AI techniques for the purposes of trademark examination as well. Australia uses a Smart Assessment Toolkit that relies on natural language processing and internal software to detect substantially similar trademarks, and the office in Singapore uses machine learning techniques to measure and suggest parameters to measure trademark distinctiveness.¹²¹ Of course, like other areas of AI applications, there are significant risks associated with automated decision-making, some of which stem from the legal and cultural risks associated with lack of transparency, unrepresentative training data, or difficulty in explainability, which we address more below.¹²² Particularly in the context of trademark law, which relies on subjective, context-dependent assessments, AI-driven technologies may be less useful in terms of evaluating distinctiveness, likelihood of confusion, and other variables that require a nuanced evaluation.¹²³

Of course, one additional consideration for the success of AI in trademark law involves the need for accurate, structured, multi-jurisdictional and comprehensive data.¹²⁴ Towards this end, scholars Anke Moerland and Conrado Freitas have distinguished between two different types of data: legal data, that involves decisions, oppositions and invalidity proceedings, and case law from various jurisdictions, in order to improve the accuracy of legal predictions; and market-based data, which includes information about consumer preferences, product variations, goodwill, product reputation, distinctiveness, etc.¹²⁵ As they note, privacy and data protection laws can impede the collection of such data, making both types of data difficult to compile accurately and comprehensively (let alone across jurisdictions), thereby posing a challenge to the efficacy of AI-driven judgments in the global trademark ecosystem.

119. Moerland & Freitas, *supra* note 6, at 15.

120. WORLD INTELL. PROP. ORG., *supra* note 89, at 8.

121. *Id.*

122. *See id.* at 12 (noting that Australia has developed an Automated Decision-Making Governance Framework and Policy); *see also* Engstrom et al., *supra* note 64, at 50–51 (noting explainability concerns, among others, in deploying AI at the USPTO).

123. Moerland & Freitas, *supra* note 6, at 16.

124. *Id.*

125. *Id.*

C. ARTIFICIAL INTELLIGENCE IN PRIVATE TRADEMARK SEARCH AND REGISTRATION

As studies have postulated, AI carries the potential to revolutionize advertising, particularly in terms of consumer recommendations, targeted advertising, market forecasting, and speech and text recognition.¹²⁶ However, AI-related issues have been largely underexamined regarding trademarks, specifically, especially where legal doctrine is concerned.¹²⁷ Just recently in late 2019, the USPTO solicited public comments about a range of issues involving AI and IP, including issues surrounding patents, authorship and copyrightability, trademark registrability, and datasets, amongst others.¹²⁸

While the vast majority of comments received focused on copyright and data-related issues, several consistent themes emerged regarding trademark protection. As Dev Gangjee has explained, the effect of AI on trademark registration will be more subtle than its impact on copyright or patent law, which has largely been driven by a threshold question of whether autonomous agents can be considered authors or inventors and whether the resulting work

126. See *The Future of Trademark Service Providers*, WORLD TRADEMARK REV., <https://www.worldtrademarkreview.com/reports/the-future-of-trademark-service-providers> (last visited Jan. 23, 2021) (portions on file with author) [hereinafter “TM Report”]; see also Interactive Advertising Bureau, *Artificial Intelligence in marketing Report*, IAB (Dec. 9, 2019), <https://www.iab.com/insights/iab-artificial-intelligence-in-marketing/>; Lee Curtis & Rachel Platts, *AI is Coming and It Will Change Trade Mark Law*, MANAGINGIP (Dec. 8, 2017), <https://www.hgf.com/media/1173564/09-13-AI.PDF> (focusing mostly on trademark law and its effect on retail, also noting how the law must adapt to AI); Lee Curtis & Rachel Platts, *Trademark Law Playing Catch-up with Artificial Intelligence?*, WIPO MAG. (June 2020), https://www.wipo.int/wipo_magazine_digital/en/2020/article_0001.html (same); Yashvardhan Rana, *Artificial Intelligence and Trademark Law in the Digital Age*, INTERNATIONAL JURIST (July 29, 2020), <https://www.nationaljurist.com/international-jurist/artificial-intelligence-and-trademark-law-digital-age#:~:text=Such%20products%20also%20enable%20a,in%20turn%20saving%20lawyers%20time> (discussing the potential effect of AI on trademark law). Recommendation systems might also arguably spark trademark liability claims if they offer competing products to a consumer, stemming from theories of initial interest confusion. Here, the jurisprudence on keyword searches can be instructive, as well as recent case law questioning the reach of initial interest confusion, suggesting that such theories of liability are unlikely to succeed in court. Gangjee, *supra* note 6, at 1–2. See, e.g., *Multi Time Mach., Inc. v. Amazon.com, Inc.*, 804 F.3d 930 (9th Cir. 2015) (noting that clear labels by Amazon in making recommendations precluded a theory of liability); *Rescuecom Corp. v. Google Inc.*, 562 F.3d 123 (2d Cir. 2009) (Google’s use of the Rescuecom trademark was a use in commerce); *Rosetta Stone v. Google*, 676 F.3d 144 (4th Cir. 2012) (overturning a grant of summary judgment for Google).

127. See Gangjee, *supra* note 6.

128. *Request for Comments on Intellectual Property Protection for Artificial Intelligence Innovation*, USPTO (Oct. 30, 2019), <https://www.federalregister.gov/documents/2019/10/30/2019-23638/request-for-comments-on-intellectual-property-protection-for-artificial-intelligence-innovation>.

product is protectable.¹²⁹ Very few comments focused on the related question of AI-created marks. At least one commentator concluded that the possibility of AI-created marks existed but emphasized that only live humans should be able to file for registration.¹³⁰

The vast majority of trademark-related comments from organizations concluded that the use of AI would improve and streamline the trademark search and registration process, noting that “[d]ecisions to proceed or not to proceed with filing a U.S. trademark application for a particular mark may be made more quickly and may be better informed if driven by a more objective risk assessment.”¹³¹ However, the increase in accuracy, at least one commentator noted, would also raise the bar for successful trademark applications, making them potentially harder to obtain but improving the overall quality of trademarks nevertheless.¹³²

A second theme was that AI could have a transformative effect on the detection of trademark infringement with its rapid search and comparison technology, aiding the USPTO in determining fraudulent applications.¹³³ At the same time, the American Bar Association (ABA) also noted the risk that AI tools and software could be used in the opposite way—to infringe the rights of other trademark owners—thus opening up questions of machine volition and liability.¹³⁴ At least one other commentator expressed a similar view, warning that while AI could be used to better detect infringement and protect trademarks, the very same technology could also be used to violate trademark

129. Gangjee, *supra* note 6, at 1.

130. See *Commentary from A-CAPP*, USPTO 1 (Dec. 16, 2019), https://www.uspto.gov/sites/default/files/documents/Jeffrey-Rojek_RFC-84-FR-58141.pdf (noting that “the creation of the trademark itself should not be allowed by AI, emphasizing role for humans in registration”).

131. ABA Letter, *supra* note 76, at 5; see also *Letter from Computer & Communications Industry Association and Internet Association to Secretary of Commerce for Intellectual Property & Director of the United States Patent and Trademark Office*, USPTO 10 (2020) (noting searches would be faster and more efficient) [hereinafter *Computer & Communications Industry Letter*]; Trevor Little, *Lower risk applications, increased refusals and a boost for infringers: the potential impact of AI on trademarks*, WORLD TRADEMARK REV. (Mar. 23, 2020), <https://www.worldtrademarkreview.com/anti-counterfeiting/lower-risk-applications-increased-refusals-and-boost-infringers-the>.

132. See generally *Letter from Obeebo, Inc. to Secretary of Commerce for Intellectual Property & Director of the United States Patent and Trademark Office*, USPTO, https://www.uspto.gov/sites/default/files/documents/Obeebo-Inc_RFC-84-FR-58141.pdf (noting that AI will raise the bar for distinctiveness, but ultimately improve trademark quality).

133. See ABA Letter, *supra* note 76, at 12 (noting that AI could aid a pixel-by-pixel comparison); *Comments from the App Association*, USPTO 5 (date goes here), <https://www.uspto.gov/initiatives/artificial-intelligence/notices-artificial-intelligence-non-patent-related> (noting that AI tools are used to detect infringement).

134. ABA Letter, *supra* note 76, at 13.

rights as well.¹³⁵ The commentary, from a center focused on anti-counterfeiting, warned that AI could be used to detect gaps in trademark protection and deceive consumers with strategically driven recommendations.¹³⁶ “At what level of prediction is there a duty to inform consumers, or b[r]and owners, about a potentially suspicious product?” the commentary asked, noting a potentially increased risk of inaccuracy from AI-driven counterfeit detection.¹³⁷ Here, if an AI tool makes an infringing recommendation, consumer harms might stem not from initial interest or point-of-sale confusion, but rather from the harm of post-sale confusion.¹³⁸

A third theme involved the consistent idea that the law did not need reforming due to the advent of AI, although many expressed a desire to avoid weakening trademark protection as a result of AI.¹³⁹ One representative view, along similar lines, expressed by the ABA and several others, involved the conclusion that AI could serve as “an appropriate supplement, but not a substitute for the human judgment of [counsel].”¹⁴⁰ Similarly, another set of commentators observed that using AI to supplement (rather than supplant) human judgment would avoid the risk that complete reliance on AI might produce an incorrect conclusion.¹⁴¹ At least one study echoed this view by

135. See *Comments from the Center for Anti-Counterfeiting and Product Protection*, USPTO (Dec. 16, 2019), https://www.uspto.gov/sites/default/files/documents/Jeffrey-Rojek_RFC-84-FR-58141.pdf.

136. *Id.* at 2–3.

137. *Id.* at 3.

138. See Trevor Little, *Lower Risk Applications, Increased Refusals and a Boost for Infringers: The Potential Impact of AI on Trademarks*, WORLD TRADEMARK REV. 2 (Mar. 23, 2020), <https://www.worldtrademarkreview.com/anti-counterfeiting/lower-risk-applications-increased-refusals-and-boost-infringers-the> (quoting commentary from the American Intellectual Property Law Association).

139. See *Comments from the App Association*, *supra* note 133, at 5 (noting a desire to avoid weakening trademark law); see also Computer & Communications Industry Letter, *supra* note 131, at 10 (noting no impact of AI on trademark law, and no need to change the law at this time).

140. ABA Letter, *supra* note 76, at 12 (noting that AI should not be used as a substitute for subjective judgment); see also *Letter from IBM Corporation to Secretary of Commerce for Intellectual Property & Director of the United States Patent and Trademark Office*, USPTO 5 (Jan. 19, 2019), <https://www.uspto.gov/initiatives/artificial-intelligence/notices-artificial-intelligence-non-patent-related> (noting that a trademark examiner will still be required to assess the evidence collected in the examination and registration process).

141. *Letter from Japan Intellectual Property Association to Secretary of Commerce for Intellectual Property & Director of the United States Patent and Trademark Office*, USPTO 2 (Jan. 8, 2020), <https://www.uspto.gov/initiatives/artificial-intelligence/notices-artificial-intelligence-non-patent-related>; see also Intellectual Property Owner’s Association 6, available at <https://www.uspto.gov/initiatives/artificial-intelligence/notices-artificial-intelligence-non-patent-related>.

noting that the subjectivity and complexity of trademark law's doctrinal tests would be difficult to replicate with an AI-driven system, since they are presently unable to reflect the nuances of these tests.¹⁴²

Yet most commentary noted, as applied specifically to trademarks, AI carries perhaps the strongest potential in areas of private search and registration.¹⁴³ More recent tightening of corporate budgets, coupled with improvements to AI technology, have streamlined the potential for AI to have a transformative effect on the process of trademark registration and litigation.¹⁴⁴ Here, AI-powered search takes a form that is much more predictive in nature, since it is primarily concerned with giving a potential registrant information about whether a preexisting registration will cause their application to be rejected. This type of search can range in complexity. At its most basic, a search engine might check to see if an application exactly matches an existing registration. More complex implementations might use AI to determine the likelihood that the USPTO would reject an application by modeling their own decision-making process. Other techniques might be most advantageous when they can be used to automate tasks like trademark search and watch results.¹⁴⁵ Since AI provides great improvements in terms of speed and accuracy, it can dramatically assist brands who aim to be the first to reach the market.¹⁴⁶

While a comprehensive view of all of the implications of AI for trademark law is beyond the scope of this article, it bears mentioning that we can envision at least five different ways in which AI-related technologies can radically alter our existing legal systems, and drive the processes of search and registration to

142. Moerland & Freitas, *supra* note 6, at 2.

143. See TM Report, *supra* note 126, at 1 (page number corresponds to excerpts on file with author).

144. See *id.* at 3 (page number corresponds to excerpts on file with author).

145. See Rob Davey, *Artificial Intelligence: A Meeting of Minds*, WORLD TRADEMARK REV. (Nov. 1, 2017), <https://www.worldtrademarkreview.com/portfolio-management/artificial-intelligence-meeting-minds>.

146. In particular, models that draw on fuzzy logic are particularly well suited for knowledge that contains elements of vagueness, like knowledge based on natural language. Anna Ronkainen describes how type-2 fuzzy logic systems are particularly appropriate for representations of second-order vagueness, especially in situations, like trademarks, where there may be a “vagueness of a concept and [an] uncertainty associated with its application.” Anna Ronkainen, *MOSONG, a Fuzzy Logic Model of Trade Mark Similarity*, in PROCEEDINGS OF THE WORKSHOP ON MODELING LEGAL CASES AND LEGAL RULES 23–25 (Adam Z. Wyner ed., 2010). In simple terms, Ronkainen writes, “traditional fuzzy logic allows us to say that John is 0.9 TALL (whatever that means), whereas with type-2 fuzzy logic we can also say that John is between 0.85 and 0.95 (0.90 +/- .05 TALL), in which the uncertainty or margin of error may stem from any source, anything from potential measurement errors to intrinsic design factors within the model.” *Id.*

be much more proactive in terms of identifying variables that can prove determinative later on.¹⁴⁷ Some examples involve the following:

1. *Search, Identification, and Suggestion*

AI carries the potential to help trademark owners search and identify potential trademarks for registration by employing AI to study a wide range of variables relevant to the search process including sight, sound, visual cues, classification of goods/services, and other trademark attributes like descriptiveness. But this can also integrate other external considerations in its analysis, like identifying geographic areas of potential growth, obstacles for trademark goodwill, other similar trademarks, or by noting attributes of other firms within the trademark ecosystem.

The same observation can easily be made for the role that AI and machine learning techniques play in the process of trademark selection.¹⁴⁸ AI can direct the trademark firm applicants to various options that are curated for them, drawing from a vast expanse of market-based data on consumer preferences, brand equity, common law variations, linguistic sophistication, natural language associations, and the like. Search and registration can also be improved using AI techniques, where machine learning can be relied upon to identify semantically similar marks.

2. *Registration and Clearance*

AI carries the potential to revolutionize the process of registration, both in terms of automating the processes of registration and in terms of identifying particular areas where there may be conflicting registrations, and even drafting initial registrations or filings and general portfolio management.¹⁴⁹ An expert notes, “[B]rand owners will be able to clear a campaign in weeks or even days, which is essential given how quickly products and services are developed and expand.”¹⁵⁰ Another expert adds, “Naming decisions will happen in real time.”¹⁵¹ As these comments suggest, not only can tools “clear” certain proposed marks for registration, but they can also register marks with automated tools.

147. See TM Report, *supra* note 126, at 4 (page number corresponds to excerpts on file with author).

148. See Moerland & Freitas, *supra* note 6, at 4 (describing how machine learning operates in the trademark context).

149. See TM Report, *supra* note 126, at 3 (“Areas where AI will dominate include searching and clearance, prosecution (at least for simple marks), renewals and possibly even oppositions.”) (page number corresponds to excerpts on file with author).

150. See *id.* at 4 (page number corresponds to excerpts on file with author).

151. *Id.* (page number corresponds to excerpts on file with author).

3. *Comparison and Determining Substantial Similarity*

AI can alter the processes of investigating substantial similarity by relying on deep learning and fuzzy logic techniques to evaluate comparisons of trademarks and product attributes. It can investigate multiple types of similarity—visual, semantic, and image—in seconds.¹⁵² By using neural network technologies, entities can process large amounts of data in order to determine semantic equivalence, providing insights into substantial similarity and trademark relatedness.¹⁵³ As Anna Ronkainen further explains:

Trademark similarity search . . . requires searching for dissimilar images as opposed to the more common approach of searching similar (or identical) images. In the latter, as long as the amount of similar images is sufficient, one could try to train a neural network-based model to catch similarities between images. For example, in order to teach the machine to differentiate between cats and dogs we should supply it with many images of cats and dogs. Unfortunately, in a trademarks database, this is obviously not the case. Moreover, catching differences between trademarks is far more complex since it is much harder to find pairs of similar trademarks, and on top of that, there is no formal definition of similar trademarks, as trademarks are considered to be similar only if they are *deceptively* similar.¹⁵⁴

As she notes, while there are some difficulties with training machines to capture these complexities, it is reasonable to consider that techniques will continue to improve in time, thereby assisting with the determination of substantial similarity.¹⁵⁵ Others have expressed similar concerns, noting that determining trademark distinctiveness, the relevant public, the proper classification of goods and services, among other elements, are so subjective that they pose challenges to the development of AI in trademark law.¹⁵⁶

152. Visual similarity involves the question of whether two trademarks are visually similar; semantic similarity involves whether the trademarks contain the same meaning and semantic content; and text similarity involves whether the actual text of the trademark is similar. Idan Mosseri et al., *How AI will Revolutionise Trademark Searches*, WORLD TRADEMARK REV. (July 2, 2019), <https://www.worldtrademarkreview.com/ip-offices/how-ai-will-revolutionise-trademark-searches>.

153. See generally TM Report, *supra* note 126 (excerpts on file with author).

154. See Ronkainen, *supra* note 146, at 23–25 (discussing the difficulties in training an AI program to catch differences between trademarks).

155. *Id.*

156. Moerland & Freitas, *supra* note 6, at 20–23.

4. *Prediction and Risk Assessment*

As with each of the other areas, the real payoff of AI lies in its ability to predict the outcomes of various trademark-related decisions—such as the litigation risk involved in proceeding with a particular trademark or product—and the market implications of making certain choices.¹⁵⁷ Risk assessments are very useful; as Gangjee notes, “[w]hile human expertise continues to assess the conflicts results lists generated by algorithms, for risk-averse commercial clients it is extremely tempting to be guided by clearly defined percentages of similarity.”¹⁵⁸ Indeed, predictive analytics can prove to be transformative in helping businesses both create and sustain a strong presence in the marketplace, predicting the outcome of filing suit, sending a cease-and-desist, articulating various claims, or deciding whether and for how much to settle. And this is just the tip of the iceberg. Imagine every aspect of a trademark claim—its probable outcome automated, calculated, predicted and ready for real-time decision-making.

Nevertheless, despite the improvements AI will provide regarding trademark registration and litigation, it is important to note that experts continue to emphasize the importance of human oversight and participation, particularly in terms of using human judgement in complex cognitive tasks, especially in the context of trademark doctrines which are highly context-specific. This is especially true in more complex cases of multi-word or slogan marks, where humans are likely to be the best at determining areas of particular strength.¹⁵⁹

5. *Brand Management*

Finally, nearly every private trademark search engine company in our study offers brand protection services in addition to their trademark search services in some capacity.¹⁶⁰ These brand protection services generally include some combination of active monitoring of U.S. and global databases, and sometimes

157. See TM Report, *supra* note 126, at 4 (page number corresponds to excerpts on file with author).

158. See Gangjee, *supra* note 6, at 13.

159. See generally TM Report, *supra* note 126 (excerpts on file with author). One example of this, experts suggest, is having a team of humans who can physically review and correct the data from national trademark registries to ensure that proprietary trademark databases have correct examples, deleting, for example, cases where the word mark does not match the image (errors which are easy for automated systems to overlook). See generally *id.*

160. See, e.g., *Quickly respond to potentially infringing trademark applications with a powerful suite of watch solutions*, COMPUMARK, <https://www.compumark.com/solutions/trademark-watching/watching> (discussing CompuMark’s trademark watching services).

tools for pursuing legal enforcement of trademark rights against potential infringers.¹⁶¹

At a later phase of search, current trademark holders might engage in a proactive process of brand management, vigilantly searching for newly registered marks that may threaten to dilute the strength of the older trademark holder's mark. Because of the huge search costs in finding potentially conflicting trademarks, trademark owners could face a daunting proposition in attempting to enforce their trademark rights themselves. This is essentially the same problem that confronts potential registrants, who must filter out the noise and recover actual conflicts, as we have previously asserted in this paper.

Here, again, as we have suggested, AI and machine learning techniques can offer mark owners a substantial advantage in brand management and enforcement. We have strong theories about why trademarks are valuable for owners and consumers; they reduce the friction created by information asymmetries and thus facilitate useful transactions.¹⁶² Brand management is important because trademark owners need to maintain the strength of their marks in order to reduce information asymmetries.¹⁶³ Moreover, brand protection is a critical service because the USPTO explicitly says that it is not responsible for trademark enforcement; it explicitly places this burden on trademark holders.¹⁶⁴

Regardless of the reason, the additional benefit that these firms provide to their clients fits into the broader story of how the private sector is able to utilize AI in a way that gets ahead of government resources, supplementing when needed. This is discussed further below.

III. A COMPARATIVE ASSESSMENT OF THE PRIVATE SECTOR IN TRADEMARK SEARCH

One of the reasons we decided to write this Article is related to another overall observation: aside from a few prominent, recent pieces,¹⁶⁵ there is not a great deal of empirical research on trademark ecosystems, especially compared to other areas of IP. Moreover, while trademark law as a field of study has been thoroughly theorized, there is little to no systematic evidence that compares the various private vendors in the process of trademark search

161. See *infra* Section III.B.2 (full descriptions of each search engine).

162. See generally Wessells, *supra* note 16.

163. See generally *id.*

164. U.S. PATENT AND TRADEMARK OFFICE, PROTECTING YOUR TRADEMARK: ENHANCING YOUR RIGHTS THROUGH FEDERAL REGISTRATION 3 (2019) (“You, as the mark owner, are solely responsible for enforcement [of your trademark].”).

165. Rai, *supra* note 3; Ronkainen, *supra* note 146.

and registration. One relatively recent study identified fewer than seventy articles involving empirical analysis of trademarks.¹⁶⁶ While some areas involved studies of the relationship between trademarks, innovation, and firm performance, the relevant law review literature is still somewhat thin.¹⁶⁷ Other empirical pieces in trademark law have focused on questions of scarcity,¹⁶⁸ the extent of trademark dilution,¹⁶⁹ or the relationship between trademarks and innovation.¹⁷⁰

The problem of firm search costs in trademark search, therefore, lies at the periphery of these various literatures, but there is very little concrete evaluation of the issue. For example, we could find only one other study that considers how different vendors use machine learning techniques in search and registration (and this one focused mostly on government tools).¹⁷¹ Computer science literature implicitly recognizes that firms face search costs in finding potential conflicts and attempts to optimize methods that reduce these costs,¹⁷² but it does not delve into the economic consequences of deploying these methods. Similarly (and conversely), economics literature implies that the USPTO plays an important gatekeeping function in ensuring adequate search quality (i.e., that potentially damaging marks are not registered),¹⁷³ but has never addressed the question of how private vendors have emerged to respond to the USPTO's own search limitations.

Like social science literature, computer science literature is largely theoretical. Authors are primarily concerned with optimizing search algorithms and engines, rather than evaluating current implementations. There do not seem to be meta-studies that comprehensively evaluate either the visual-based search engines or text-based ones, which suggests that there are avenues for

166. See generally Shukhrat Nasirov, *The Use of Trademarks in Empirical Research: Towards an Integrated Framework* (Nov. 20, 2018) (unpublished manuscript), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3296064.

167. See *id.*

168. See, e.g., Beebe & Fromer, *supra* note 56, at 947.

169. Paul J. Heald & Robert Brauneis, *The Myth of Buick Aspirin: An Empirical Study of Trademark Dilution by Product and Trade Names*, 32 CARDOZO L. REV. 2533, 2574–75 (2011).

170. Nasirov, *supra* note 166.

171. See generally Moerland & Freitas, *supra* note 6.

172. See Fatahiyah Mohd Anuar, Rossitza Setchi & Yu-Kun Lai, *A Conceptual Model of Trademark Retrieval Based on Conceptual Similarity*, 22 PROCEDIA COMPUT. SCI. 450, 451 (2013) (“[I]n the Internet age, it is even more important to have efficient mechanisms for protecting trademarks and tools for detecting possible cases of infringement” to motivate the importance of developing their trademark conceptual similarity model.).

173. See generally Landes & Posner, *supra* note 42 (framing the economics of trademark law as being grounded in the economics of property and tort law). They argue that trademark creates a property right, and trademark litigation is a branch of tort law. *Id.* Since the USPTO grants the mark, it effectively is responsible for determining who gets a property right.

future research. The computer science literature is more directly concerned with the efficacy of search algorithms, focusing on more complex problems than standard text-based retrieval. In general, the major problem currently being tackled in the computer science literature is improving trademark retrieval based on visual similarity, instead of spelling or phonetic similarity. Various papers explore different similarity metrics that determine whether proposed design marks are similar to current marks.¹⁷⁴ One 1987 study looked at a system that could return trademark searches based on phonetic similarity, and this application is closer to the technology deployed by the search engines in our study.¹⁷⁵

In terms of literature that directly looks at the applications of text-based trademark retrieval, there are very few. Anke Moerland and Conrado Freitas conducted a study that combined qualitative methods with small-n search tests on the United Nations, European Union's, Australia's, and Singapore's public-facing trademark search engines. Moerland and Freitas note that each of these offices currently uses AI methods to power their search algorithms, whereas the USPTO is still developing and testing AI tools to identify grounds for refusals to register trademark applications.¹⁷⁶ While the study was highly illuminating, it was driven more towards examining government use of AI-related tools (as opposed to studying the private market for assessing trademark search). The study also assessed the functionality of these tools in identifying and comparing visual and conceptual similarity, descriptiveness, morality, and classifications of goods between marks.¹⁷⁷

In terms of other studies, one 1999 paper examined the potential future of patent and trademark librarians in a time when databases were becoming more

174. See generally Anuar, *supra* note 172; Anil K. Jain & Aditya Vailaya, *Shape-Based Retrieval: A Case Study with Trademark Image Databases*, 31 PATTERN RECOGNITION 1369 (1998); Gianluigi Ciocca & Raimondo Schettini, *Similarity Retrieval of Trademark Images*, in PROCEEDINGS 10TH INTERNATIONAL CONFERENCE ON IMAGE ANALYSIS AND PROCESSING 915 (Bob Werner ed., 1999).

175. J. Howard Bryant, *USPTO's Automated Trademark Search System*, 9 WORLD PAT. INFO. 5 (1987).

176. Moerland & Freitas, *supra* note 6, at 6 (discussing methodology). The qualitative semi-structured survey was sent to fourteen stakeholders, including TrademarkNow and the USPTO, which we include in our assessment. The search engine tests involved searching marks related to Apple Inc., using both its logo and the word "apple" to see if each engine flagged potential issues with the search. Specifically, they measure whether searching "apple" raises conceptually similar marks across all trademark classes and within specific classes like "fruits" and "software/hardware." As we detail below, we conduct similar tests on private sector trademark search engines using over a hundred search terms.

177. *Id.* at 7–8.

common.¹⁷⁸ The author ultimately concluded that librarians would still have a place in helping users navigate these databases, in part because of the huge volume of applications.¹⁷⁹ However, there are no retrospective studies that indicate how this prediction bore out, or how the growth of patent and trademark databases have altered applications or research.

In one related study, Lisa Larrimote Ouellette directly tackles the question of the PTO using search engines in trademark applications. Her main argument is that Google is an underexplored tool in assessing the distinctiveness of a trademark, where distinctiveness is “the extent to which consumers view a mark as identifying a particular source.”¹⁸⁰ She argues that Google, with its complex algorithm and public results, provides an easy way for the PTO to assess distinctiveness in cases of infringement. To prove her argument, she conducts an empirical experiment where she used trademarks that were disputed for trademark infringement and searched them through Google.

The basic test for whether a trademark was distinctive in this framework came down to whether it was findable in Google. If a mark was distinctive or commercially popular, then it would dominate the top ad results. If the mark was likely to be confused with another mark, there would be overlapping results between searches for those marks. Essentially, she argues, valuably, that Google can take a lot of the guesswork out of determining whether consumers would be able to discern one mark from another, and that this potential role has been underexplored in infringement cases.

Ouellette’s insightful study foregrounds the role of private companies in facilitating search and comparison and points out the potentially powerful (and troubling) role of “algorithmic authority” in trademark law.¹⁸¹ This study differs from ours primarily in that we are focused on search engines that specialize in trademark search, prior to registration, and we approach the problem from the perspective of a trademark applicant, as opposed to a traditional consumer. Ouellette’s solution is mainly relevant for courts deciding a trademark infringement case, whereas we are examining trademark applications well before they would get to that stage of the legal process.¹⁸²

178. Julia Crawford, *Obsolescence or Opportunity? Patent & Trademark Librarians in the Internet Age*, 21 WORLD PAT. INFO. 267 (1999).

179. *Id.*

180. Lisa Larrimore Ouellette, *The Google Shortcut to Trademark Law*, 102 CALIF. L. REV. 351 (2014).

181. *See id.* at 368 (citing Clay Shirky’s observations) (citation omitted).

182. Moreover, Google differs from trademark search engines in that Google’s PageRank algorithm relies on a calculation of how different webpages point to each other. *See How Google’s Algorithm Rules the World*, WIRED (Feb. 22, 2010), <https://www.wired.com/2010/02>

Trademark search engines differ in that they are more akin to querying a database, modeling good results, and returning those results to the user at a much earlier stage involving trademark selection and application. Nevertheless, her observations about the potential role of algorithms in assisting legal determinations are salient to our study, since many private search engines raise similar questions about efficacy, impact, and accuracy.

Other than the ones mentioned above, we could find only two other papers that conducted empirical tests on actual trademark search engines. Anna Ronkainen conducted a study of thirty thousand trademarks conflicts on the TrademarkNow platform.¹⁸³ She specifically models a trademark similarity algorithm developed by Onomatics, a Finland-based legal technology firm.¹⁸⁴ The Onomatics algorithm was used to power TrademarkNow's Namecheck product, and she argues it is especially good at incorporating the role of goods and services in its similarity calculation.¹⁸⁵ Her basic results showed that the algorithmic approach recovered marks with precision of about 80% and recall of about 94.9%.¹⁸⁶ A paper entitled "Trademark Search Tools" put forward by ipPerformance in 2011 is the only one that directly looks at leading trademark search vendors and does an apples-to-apples comparison of them.¹⁸⁷

A. TRADEMARK SEARCH AND REGISTRATION PROCEDURES

As discussed earlier, potential trademark registrants file their trademarks by filing an application with the USPTO. The USPTO advises that registrants should first determine whether a trademark is the appropriate protection (instead of a patent or copyright), and then details several steps for registration.¹⁸⁸ In particular, it says that registrants should select their mark, choose a format, identify whether it is a good or service mark, search for potential conflicts, and choose a filing basis.¹⁸⁹

/ff_google_algorithm/ [https://web.archive.org/web/20140412235725/http://www.wired.com/2010/02/ff_google_algorithm/].

183. Anna Ronkainen, *Intelligent Trademark Analysis: Experiments in Large-Scale Evaluation of Real-World Legal AI*, in PROCEEDINGS OF THE 14TH INTERNATIONAL CONFERENCE ON ARTIFICIAL INTELLIGENCE AND LAW 227 (Ass'n for Computing Mach. ed., 2013).

184. *Id.*

185. The trademark similarity algorithm is derived from the MOSONG prototype, which is a model of vagueness and uncertainty in legal text. *Id.* at 2.

186. *Id.*; see also *infra* Section III.B.7 (formal descriptions of terms used here).

187. IPPERFORMANCE GRP., TRADEMARK SEARCH TOOLS: ANALYSIS PAPER 7 (2011), https://www.markify.com/pdf/Trademark_Search_Tools_Analysis_Paper-P2a.pdf.

188. See *Trademark Basics*, USPTO, <https://www.uspto.gov/trademarks-getting-started/trademark-basics>.

189. *Id.* ("Other initial considerations").

Once these steps are complete, the registrant fills out an application, specifying both the mark and the class of goods upon which it will appear, and then monitors its status for USPTO approval. Crucially, each mark must be categorized as either a good or service, and the applicant must select the number of classes.¹⁹⁰ Each class costs an additional \$225–275 depending on the specific applicable fee schedule.¹⁹¹ It may be necessary to communicate with a USPTO examining attorney to talk through any potential issues or objections before getting an official approval or denial. If approved, the registrant is still responsible for enforcing the trademark.¹⁹²

As we well know, in the conventional case, trademark registrants will want to avoid the costs associated with filing a rejected trademark application. To avoid incurring these costs, they turn to trademark search engines to identify potential conflicts in advance and to make appropriate changes prior to filing a trademark application. The first step is generally to check the USPTO's TESS.¹⁹³ However, while TESS can return existing trademarks that are similar to the search term, as we have suggested, it is not totally effective. The USPTO itself recommends consulting an attorney before filing an application as it cannot guarantee that its results will be exhaustive.¹⁹⁴

In the typical use case for a trademark search engine, a potential registrant, or their attorney, searches a potential mark and then sorts through the returned results. Firms may employ attorneys to conduct a trademark search, and, consequently, attorneys turn to trademark search engines to assist with this process. Attorneys need to be exceptionally careful when advising their clients,

190. See Engstrom et al, *supra* note 64, at 46–47 (description of the trademark process before the USPTO).

191. For more details on the trademark application form, see *Trademark Initial Application Form*, USPTO, <https://www.uspto.gov/trademarks-application-process/filing-online/initial-application-forms#Chart%20Application%20requirements> (last visited July 28, 2019). For details on the fee schedule, see *USPTO Fee Schedule*, USPTO, <https://www.uspto.gov/learning-and-resources/fees-and-payment/uspto-fee-schedule#TM%20Process%20Fee> (last visited April 13, 2020).

192. *Trademark Process*, USPTO, <https://www.uspto.gov/trademarks-getting-started/trademark-process#step3> (last visited July 28, 2019).

193. *Search Trademark Database*, USPTO, <https://www.uspto.gov/trademarks-application-process/search-trademark-database> (last visited July 16, 2018).

194. *Id.* Specifically, the website advises:

[D]eciding what to search for and interpreting your results can be complicated. There are many factors to consider in determining **likelihood of confusion**. We can't advise you on how to do a clearance search for your mark, do one for you, or interpret your search results. Therefore, we strongly encourage you to **hire a U.S.-licensed attorney** who specializes in trademark law to guide you throughout the application process.

Id. (emphasis added).

and, therefore, trademark search engines are likely optimized in a way to ensure that attorneys can trust their results as being definitive. Routinely not returning an accurate result for a potential conflict could cause attorneys to shift business away from one search engine toward another.

Here, as discussed above, private vendors have emerged to assist attorneys and applicants with their own search processes. These trademark search engines all use some form of search algorithm to power their results, although each of them utilizes different methods of integrating data and machine learning into their analytical performance. Again, these can vary in complexity and may be geared toward different audiences. The broad takeaway is that they each represent a means of helping registrants navigate a complicated search problem by reducing search costs for marks through recent advancements in technology. By giving applicants the ability to go beyond what TESS or a library search can provide, they potentially reduce search costs considerably.

The core type of search that each search engine provides is the “knockout search.” A knockout search is essentially a trademark search that intends to return marks that are likely to be cited in a 2(d) “likelihood of confusion” rejection for a new trademark application.¹⁹⁵ This category is what we focus most of our empirical analysis on because it is the one point of common ground between all of the search engines in our study. Within the knockout search, there are still some ways that different search engines distinguish themselves. Some may simply reference the USPTO’s own TESS search engine,¹⁹⁶ while others combine that data with their own methods.¹⁹⁷ Still others will attach likelihoods for risk scores, which requires a more algorithmic approach than simply checking against TESS.¹⁹⁸

195. *What is a Trademark Knockout Search?*, PAT. TRADEMARK BLOG, <http://www.patenttrademarkblog.com/trademark-knockout-search/> (last visited July 28, 2019). According to the Trademark Manual of Examining Procedure, likelihood of confusion refers to a mark that, “as used on or in connection with the specified goods or services, so resembles a *registered* mark as to be likely to cause confusion.” TMEP § 1207.01, available at <https://tmap.uspto.gov/RDMS/TMEP/current#/current/TMEP-1200d1e5044.html>.

196. Trademarkia’s free service does this, for example. *See* TRADEMARKIA, <https://www.trademarkia.com/> (last visited July 28, 2019).

197. *See, e.g.*, TRADEMARKNOW, <https://www.trademarknow.com/> (last visited July 28, 2019) (optimizing for speed by prioritizing returning “exact matches”). In its ExaMatch (<https://www.trademarknow.com/products/examatch>) page, it includes a search engine to search the USPTO and E.U. databases.

198. Both Markify and CSC provide likelihood measures with their results. *See* MARKIFY, <https://www.markify.com> (last visited July 28, 2019); *see also* CORPORATION SERVICE COMPANY (CSC), <https://www.cscglobal.com/global/web/csc//trademark-searching.html> (last visited July 28, 2019); *infra* Figure 5.

Different search engines differentiate their core products, so making comparisons between them necessarily simplifies the typical use case for each one. Different search engines will provide different metrics, and there are a few other considerations as well.¹⁹⁹ Many of the search engines in our study distinguish themselves by offering a “comprehensive search” of some sort.²⁰⁰ These services can vary considerably between different search engines. One major consideration is whether a comprehensive search involves automation or human review. Some comprehensive search tools will automatically generate detailed reports, whereas others have human beings thoroughly investigate a potential mark.

In sum, because of the diversity in trademark search products, evaluating their performance can be tricky. Trademarks have several different elements, and there are multiple ways that a trademark application can be “confusingly similar.” Moreover, identifying a “confusingly similar” registration involves some judgment as well because different search engines could return noisier results than others, even if the “correct” answer is present in all of them. Namely, a trademark application can be similar to an existing one in its visuals, phonetics, concept, or spelling.²⁰¹ To address this issue, Idan Mosseri and colleagues created “TradeMarker” software, which conducts a variety of independent searches, developing metrics of automated content similarity, image/pixel text similarity, and manual content similarity.²⁰² They construct individual similarity measures for each of these categories, and then combine

199. For instance, whether a conflicting mark is “live” or “dead” is relevant as a dead mark cannot be cited as a reason to reject a proposed mark. *Searching Marks in USPTO Database*, USPTO, <https://www.uspto.gov/trademarks-getting-started/trademark-basics/searching-marks-uspto-database> (last visited July 28, 2019).

200. Trademarkia explicitly talks about a comprehensive search, while others like Corsearch offer a “trademark screening platform.” See TRADEMARKIA, *supra* note 196; *Trademark Screening*, CORSEARCH, INC., <https://www.corsearch.com/our-products/trademark-screening/> (last visited July 28, 2019).

201. See *Possible Grounds for Refusal of a Mark*, USPTO, <https://www.uspto.gov/trademark/additional-guidance-and-resources/possible-grounds-refusal-mark> (last visited July 28, 2019).

202. Idan Mosseri, Matan Rusanovsky & Gal Oren, *TradeMarker – Artificial Intelligence Based Trademarks Similarity Search Engine*, SPRINGER NATURE SWITZ. 97 (2019), available at https://www.researchgate.net/publication/334352698_TradeMarker_-_Artificial_Intelligence_Based_Trademarks_Similarity_Search_Engine/link/5d2865cd458515c11e27b220/download; see also Tim Lince, *How AI will revolutionize trademark searches*, WORLD TRADEMARK REV. (July 2, 2019), <https://www.worldtrademarkreview.com/ip-offices/how-ai-will-revolutionise-trademark-searches> (highlighting guest analysis provided by TradeMarker that combines visual, semantic/content, and text similarity).

each of these measures for an “overall similarity” score.²⁰³ This mixture is useful because it avoids situations where two marks are unlikely to be considered “confusingly similar,” even if they share some aspect of their marks. For example, Target and Vodafone have very similar logos, but do not share text, conceptual, or spelling similarities and therefore would not have a high combined similarity score.

It is also worth noting that each trademark search engine firm offers services beyond just search. Many search engines offer active trademark screening, which takes a client’s existing marks and checks to see if potential conflicting marks have been applied for or registered.²⁰⁴ Again, the USPTO does not take responsibility for enforcing trademarks against potential infringers,²⁰⁵ and therefore likely created a market for these technologies. This sort of service gives companies the ability to engage in brand management. Brand management is at the core of why trademark law exists and is of high importance to firms, and the searches involved with these activities likely mirror the core technologies powering the core search engine functionality.²⁰⁶ Below, we outline some of the major characteristics of the trademark search engines we studied.

B. A COMPARISON OF TRADEMARK SEARCH ENGINES

1. *Public Search Engines*

a) USPTO

The USPTO offers TESS to search existing trademarks for a potential conflict. TESS allows users to search marks that have been both registered and applied for, but it does not automatically flag conflicts on its own. Instead, the USPTO suggests that users supplement a TESS search by consulting an attorney or using a trademark search firm.²⁰⁷ TESS further offers a few different options for search inclusiveness, depending on the user’s sophistication. Its basic search function does a simple search for word

203. See Mosseri et al., *supra* note 202 (“This separation enables us to benefit from the advantages of each aspect, as opposed to combining them into one similarity aspect and diminishing the significance of each one of them.”).

204. For example, both Markify and Corsearch offer these services; see descriptions of each search engine below.

205. U.S. PAT. & TRADEMARK OFFICE, *supra* note 164.

206. See generally Landes & Posner, *supra* note 42 (discussing the economics of trademark’s signaling quality to consumers).

207. *Search Trademark Database*, *supra* note 193 (see “Trademark Searching” and “Hiring an Attorney”).

matches, whereas its more advanced engines use design mark codes and other information to construct results.

TESS was launched in 2000, making it one of the oldest systems in our study.²⁰⁸ At the time it was launched, the USPTO explained that TESS used the same search engine and database that its own examiners use.²⁰⁹ However, few details are available about the exact search algorithm. One main disadvantage of TESS is that it seems to have relatively few computational resources, as only a fixed number of people may search at once and it requests that users log out to release resources to others in the queue.²¹⁰ Previously, the USPTO offered a different free search service since 1998, but TESS ultimately replaced it.

Importantly, TESS also draws from the U.S. government's trademarks dataset.²¹¹ The trademark case files dataset²¹² contains information about over eight million trademarks and is the authoritative source for existing and previous trademarks in the United States. The advantage of TESS is that it draws directly upon this dataset, and consequently uses it to generate its own search results.

Although its underlying search algorithm and use of AI is unclear, TESS does have a number of useful features for potential registrants. It provides serial numbers, registration numbers, and whether a conflicting mark is live or dead, like shown in Figure 2. Some ordering occurs as exact matches tend to appear near the top of the search results, but this exact mechanism has not been verified.

208. Press Release, *USPTO Introduces New Trademark Electronic Search System*, USPTO (Feb. 29, 2000), <https://www.uspto.gov/about-us/news-updates/uspto-introduces-new-trademark-electronic-search-system>.

209. *Id.*

210. *Trademark Search: Beginners Guide to Everything to Know*, UPCOUNSEL, INC., <https://www.upcounsel.com/trademark-search> (last visited July 28, 2019).

211. Stuart J.H. Graham, Galen Hancock, Alan C. Marco & Amanda Myers, *The USPTO Trademark Case Files Dataset: Descriptions, Lessons, and Insights*, 22 J. ECON. & MGMT. STRATEGY 669 (2013); *see also* Trademark Electronic Search System (TESS), USPTO, <http://tmsearch.uspto.gov/bin/gate.exe?f=tess&state=4806:pvkuk8.1.1> (last visited Jan. 23, 2021) (“This search engine allows you to search the USPTO's database of registered trademarks and prior pending applications to find marks that may prevent registration due to a likelihood of confusion refusal.”).

212. *See Trademark Case Files Dataset*, USPTO, <https://www.uspto.gov/learning-and-resources/electronic-data-products/trademark-case-files-dataset-0> (last visited July 28, 2019).

Figure 2: TESS Search Results

[Home](#) [Site Index](#) [Search](#) [FAQ](#) [Glossary](#) [Guides](#) [Contacts](#) [Feedback](#) [eBiz Alerts](#) [News Help](#)

[Trademarks > Trademark Electronic Search System \(TESS\)](#)

TESS was last updated on Mon Aug 5 04:51:02 EDT 2019

Logout Please logout when you are done to release system resources allocated for you.

Shift Left Alt: OR Jump to record: **62 Records(s) found (This page: 1 ~ 50)**

Refine Search: (SERIES 1)(COMB) Submit

Current Search: 31: (SERIES 1)(COMB) docs: 62 doc: 326

Serial Number	Reg. Number	Word Mark	Check Status
1	8814046	EX SERIES ZONE 1	TSDR LIVE
2	0793989	ECONO GATE SERIES 1	TSDR DEAD
3	0793978	TOUGHY UPDATE SERIES 1	TSDR LIVE
4	0793967	ATLANTS G-GUTTER CANOPY SERIES 1	TSDR LIVE
5	0749036	KITCHEN ESSENTIALS CLUB SERIES 1	TSDR DEAD
6	0735448	EVERY PRIZE ONLY 5 A. 1 877 FAUCION L/2 THE UNITED ES AME \$5 ONLY FIVE DOLLARS A PRIZE ONLY FIVE DOLLARS A PRIZE EVERONES FAVORITE FIVE DOLLAR BILL 877 7A BILL FINNF	TSDR LIVE
7	0718608	EVERONES FAVORITE \$5 BILL SERIES 2003 OORAH AUCTIONS HOME OF THE \$5 AUCTION FW 48 5 FIVE LINCOLN DOLLARS 5	TSDR DEAD
8	0709692	NETAS 1-SERIES	TSDR LIVE
9	0644024	LEGACY SERIES	TSDR DEAD
10	0617354	CHAMPIONS OF EDUCATION 1 MWP SERIES	TSDR LIVE
11	0610225	1 MWP SERIES CHAMPIONS OF EDUCATION	TSDR DEAD
12	0609095	100000 1000000 FEDERAL RESERVE BILL J0051201021 THE UNITED SMOKERS OF AMERICA MILLION MONEY 2 B BLOW J0051201021 THIS WRAP IS INENDED FOR SMOOKING PURPOSES ONLY THIS WRAP IS NOT A FORM OF LEGAL TENDER MONEY 2 BLOW D 4 TREASURER OF THE UNITED STATES SERIES 2013 SECRETARY OF THE TREASURY 1,000,000	TSDR DEAD
13	0597205	SQUATTERS SECRET STASH 1-OF SERIES	TSDR LIVE
14	13203800	SYNERTRAL LINK FIXED INDEX ANNUITY - SERIES 1	TSDR LIVE
15	13203738	DR LOUIS SO NATURAL BUSINESS FIRST AND ONE BY ROYAL SO NATURAL ONE INTERNATIONAL MEDICAL PHARMACEUTICAL CHEMICAL BIOLOGICAL NUTRITION & MET LABORATORY & COMPANY PRODUCTS SO NATURAL DR. LOUIS COSMETICS SO NATURAL ONE DR. LOUIS PARAPHARMAS & CREAMS PRODUCT DESCRIPTION SERIES NUMBER 10TH ANNIVERSARY 2006 - 2016 NEW ROYAL TRADEMARK DIVISIONS LONDON BRIDGE LONDON ENGLAND UK; DIVISION 1 THAMES RIVER BRIDGE (AMTRAV) NEW LONDON CONNECTOUT USA; DIVISION 2 RUSSIA SHERMERTY CASTLE LEHNINGRAD DIVISION 3 KINIGAN MOUNTAINS DIVISION 4 BRIDGE OF WIER (CROSSING POINT FOR THE RIVER GRAYFE) SCOTLAND UK; DIVISION 5 GRAND CANYON SYNERALINK COLOMBO CHITPAU PROTEOZOIC & PALEOZOIC STRAIN USA; DIVISION 6 RIVER RYNE WYRE SWITZERLAND LICHTENSTEIN USHER BERNIED FRANCE IRELANDS BATTLESHIP (AIR CRAFT CARRIER) A GLORY OF WORLD WAR II SHE IS NOW A MUSEUM SHIP AT PEARL HARBOR; DIVISION 7 THE GREAT WALL OF CHINA CHINA; DIVISION 8 HIMALAYA MYSTERAL OVER THE WORLD DIVISION 12 CITY LAS VEGAS NEVADA USA DIVISION 13 DONAVUS RIVER EUROPE DIVISION 14 CITY OF LONDON UK; DIVISION 15 CANADA DIVISION 16 SEAS FISHING BOAT 17 GAYXOER THE SWENBO 18 HOLLANDS COAST GUARDS VESSEL 19 THE BARBANTON OF LONDON ENGLAND THE LIVE FISH MARKET 20 THE WATER LAKES RIVERS 1949 - 1982 SCIENTIFIC & ART WORK WITH TOP CHEMISTS & PHARMACISTS SINCE 1989 & CONTINUE BUSINESS SINCE 2008 & CONTINUE S.M DR. S.M VAS 1. LOUIS	TSDR LIVE

2. *Private Search Engines*

a) Corsearch

Corsearch is a relative newcomer to the trademark AI space, having become its own independent company in 2017.²¹³ It is headquartered in New York City and mainly serves corporate customers, according to Crunchbase.²¹⁴ Corsearch is a “brand management” service that offers a range of tools to serve this end. These tools are largely powered by AI and search optimization techniques, and much of Corsearch’s value-add seems to be in speed, ease-of-use, and comprehensiveness.²¹⁵

The company offers an array of IP services. Namely, it offers trademark screening, trademark searching, trademark watching, online brand protection, and domain name services.²¹⁶ Basically, Corsearch provides a suite of tools for brand management, broadly construed. Its tools allow a user to screen for potential conflicts before filing, search globally once they are ready to do an exhaustive search, watch for potential new conflicts after the mark has been registered, and take legal action against potential infringers.²¹⁷ Thus, it creates a complete trademark workflow for potential registrants.

In our study, we focus on Corsearch’s “trademark screening” product. The trademark screening engine is a dashboard that provides search results for queries, along with some additional services like visualization, document creation, etc. We focus on trademark screening because it is the closest equivalent to the main trademark search product offered by all of the engines in our study.

That being said, Corsearch claims to distinguish itself from its closest competitors in a number of ways. According to its webpage, its main value lies in its “phonetic search engine.” The phonetic search engine allows a user to see results that include phonetic, spelling, and plural variations. Theoretically, this should allow the engine to cover idiosyncratic spellings, and therefore help a client do an exhaustive search for any potential conflicts. Later, we explore phonetic matches as this is one of the common ways a mark application can

213. See *Corsearch Inc.*, BLOOMBERG L.P., <https://www.bloomberg.com/profile/company/1632077D:US> (last visited July 28, 2019).

214. *Corsearch*, CRUNCHBASE, INC., <https://www.crunchbase.com/organization/corsearch#section-lists-featuring-this-company> (last visited July 28, 2019).

215. *Id.*

216. *Our Solutions*, CORSEARCH, INC., <https://corsearch.com/our-products/products-overview/> (last visited July 28, 2019).

217. *Id.*

be rejected,²¹⁸ and Corsearch plausibly has a comparative advantage in these sorts of searches.

Recently, Corsearch acquired Principium to bolster its own trademark watching services.²¹⁹ In addition to its other recent acquisitions, Corsearch is building its portfolio to ensure that it can compete on every aspect of brand management.²²⁰

Figure 3: Corsearch's Search Terminal²²¹




Figure 4: Corsearch Example Results

Trademark	International Class	Goods/Services	Status	Owner/Company
"PERFECT 1-2-3 SERIES" (Stylized)			Cancelled November 3, 1992	HELBER'S DIAMOND SHOPS, INC.
1 MILLION PLUS SERIES			Abandoned April 1, 2008	VECTOR PRODUCTS, INC.
1 MILLION SERIES			Abandoned July 27, 2007	VECTOR PRODUCTS, INC.
1 MVP-SERIES CHAMPIONS OF EDUCATION and Design			Abandoned December 1, 2014	MVP SERIES, L.L.C.
1 SERIES			Registered June 22, 2013	BAYERISCHE MOTOREN WERKE AKTIENGESELLSCHAFT
1 SERIES - BPS			Registered March 1, 2013	COMMERCIAL METALS COMPANY
1,000,000 + SERIES			Abandoned August 6, 2007	VECTOR PRODUCTS, INC.
1,000,000 POWER SERIES			Abandoned April 3, 2009	VECTOR PRODUCTS, INC.
1,000,000 SERIES			Abandoned July 31, 2007	VECTOR PRODUCTS, INC.
1,500,000 + SERIES			Abandoned May 3, 2007	VECTOR PRODUCTS, INC.
1,500,000 POWER SERIES			Abandoned April 2, 2008	VECTOR PRODUCTS, INC.
1,500,000 SERIES			Abandoned July 19, 2007	VECTOR PRODUCTS, INC.
1.5 MILLION PLUS SERIES			Abandoned April 1, 2008	VECTOR PRODUCTS, INC.
1.5 MILLION SERIES			Abandoned July 27, 2007	VECTOR PRODUCTS, INC.
100 DEAD PRESIDENTS RESERVE NOTE TX 9560207G Q2 DEAD PRESIDENTS PAPER TO			Registered October 17, 2017	LARA, RAFAEL

218. See *Possible Grounds for Refusal of a Mark*, USPTO, <https://www.uspto.gov/trademark/additional-guidance-and-resources/possible-grounds-refusal-mark> (last visited Jan. 23, 2021); see also Beebe & Fromer, *supra* note 56, at 1039 (discussing how the FDA also uses phonetic similarity to determine whether drug names are confusingly similar to one another).

219. See *Corsearch Acquires Principium Trademark Watch and Domain Services Businesses*, BUSINESSWIRE (May 17, 2019, 4:00 AM), <https://www.businesswire.com/news/home/20190517005089/en/Corsearch-Acquires-Principium-Trademark-Watch-Domain-Services>.

220. See *id.*

221. Note that it includes several language-based search parameters including phonetic search.

b) Markify

Markify was founded in 2009,²²² and is exclusively specialized in trademark searches and brand management. Markify is headquartered in Sweden and provides global services that allow clients to search and manage trademarks across numerous jurisdictions.²²³ In 2017, LegalZoom, an American legal technology company, partnered with Markify to power its own trademark and monitoring services.²²⁴ LegalZoom specializes in providing legal help to small businesses and other entities.²²⁵ One of these services is trademark registration, and LegalZoom provides a trademark search as part of its process.²²⁶ Because of LegalZoom's dominance in the U.S. market, its partnership is a key part of Markify's portfolio.²²⁷

Markify provides several services as part of its general brand management offerings.²²⁸ These include its Comprehensive Search, ProSearch, trademark watch, domain name watch, and an API.²²⁹ The ProSearch search feature is the closest equivalent to other search engines in our study, and thus we focus on this product.²³⁰ The trademark watching service actively checks international trademark databases and provides weekly reports about potential conflicts.²³¹ Similarly, the domain name watch looks for confusingly similar domain names.²³²

Markify's services are powered by its own trademark similarity search algorithm. The company argues that it distinguishes itself by developing its algorithm from a statistical perspective, so that users can prioritize search results more easily.²³³ This approach quite explicitly leverages artificial

222. See *Markify*, CRUNCHBASE, <https://www.crunchbase.com/organization/markify> (last visited Jan. 23, 2021). Please see our first footnote, noting that Markify provided funding for this study.

223. *Id.*

224. See *LegalZoom Selects Markify as Trademark Search and Monitoring Provider*, BUSINESSWIRE (Nov. 07, 2017), <https://www.businesswire.com/news/home/20171107005509/en/LegalZoom-Selects-Markify-Trademark-Search-Monitoring-Provider>.

225. *Id.*

226. Joe Runge, *Why Do I Need to Conduct a Trademark Search?*, LEGALZOOM.COM, INC., <https://www.legalzoom.com/articles/why-do-i-need-to-conduct-a-trademark-search> (last visited July 28, 2019).

227. BUSINESSWIRE, *supra* note 224.

228. See *Products & Pricing*, MARKIFY, <https://www.markify.com/> (last visited July 28, 2019).

229. *Id.*

230. *Id.*

231. *Id.*

232. *Id.*

233. It says, "The trademark search algorithm was developed by a team of mathematicians, linguists and computer scientists. It was built on a statistical analysis of more

intelligence, statistical analysis, and big data to transform trademark search.²³⁴ Markify’s central goal is to return as many potential conflicts as possible, but to also filter out as much of the “noise” as possible.²³⁵ Noise in this case would be search results that do not actually present a conflict or are not plausible 2(d) violations.

Figure 5: Markify Search Results²³⁶

The screenshot shows the Markify search results for 'CLOSET ENVY'. The interface includes a sidebar with navigation options like 'Comprehensive reports', 'Trademark watch', 'Domain watch', and 'My account'. The main content area is divided into several sections: 'Ordered search report', 'Result summary', 'Risk analysis', and 'Created reports'. Below these is a 'Filters' section and a large table of search results.

Ordered search report
 Search: CLOSET ENVY
 Order date: July 9, 2019
 Classes: 3, 9, 10, 11, 25, 41, 42, 43
 Similarity level: Comprehensive
 Search package: US
 Order #: 1064221

Result summary
 US federal: 922
 Translations: 0
 Other classes: 37
 Dead marks: 54

Risk analysis
 Very high risk: 1
 High risk: 395
 Medium risk: 66
 Low risk: 460

Created reports
 No customized report

Filters
 Goods & services filter: No filter
 General filters: No filter

Table of Results:

Risk level	Select	Add	Add full record	Row nr.	Comment	Trademark	Logotype	Classes	Database	Owner	Goods & Services	Status	Dates	Serial
1	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1		CLOSET ENVY	CLOSET ENVY	43	US (USPTO)	Marriott International, Inc.	43 - Hotel Services	Registered	AD: May 28, 2014 RD: Jun 23, 2015	80205232 RNL 4756955
2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2		ENCLOSE	ENCLOSE	10	US (USPTO)	VITALTEC INTERNATIONAL, INC.	10 - Surgical instruments and...	Registered	AD: Jul 5, 2000 RD: Oct 15, 2002	76089940 RNL 2636533
2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	3		ENDO CLOSE	ENDO CLOSE	10	US (USPTO)	COVIDEN LP	10 - medical devices; namely...	Registered	AD: Jun 28, 1992 RD: Feb 28, 1995	ARI 74389509 RNL 1881350
2	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	4		KLOUT ENVY	Klout Envy	25, 35	US (USPTO)	Klout Envy	25 - Clothing; namely, hats, shirts...	Filed	AD: Aug 23, 2018	88887042
2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	5		ENVIED CLOTHING CO.	Envied Clothing Co.	25	US (USPTO)	Russell, Tasha	25 - T-shirts; T-shirts for adults...	Registered	AD: Jul 9, 2017 RD: Sep 11, 2018	87500997 RNL 5952897
2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	6		ENVI	ENVI	9	US (USPTO)	HARRIS GEOSPATIAL SOLUTIONS, INC.	9 - computer software for analysis...	Registered	AD: Sep 21, 1993 RD: Jul 4, 1995	ARI 74438227 RNL 1902440

than 8[,]000 actual cases where a government official had ruled that two trademarks were confusingly similar. The trademark search technology is constantly upgraded and adapted to new markets.” *About Markify*, MARKIFY, <https://www.markify.com/about.html> (last visited July 28, 2019); see also *Big Promises, Big Data*, WORLD INTELL. PROP. REV. (May 21, 2019), <https://www.worldipreview.com/contributed-article/big-promises-big-data> (noting Markify’s role in harnessing big data).

234. *Id.*

235. See *Get a real comprehensive trademark watch service*, MARKIFY, <https://www.markify.com/services/trademark-watch.html> (last visited Jan. 23, 2021) (discussing its “signal-to-noise ratio”).

236. Note that in addition to the mark name, Markify returns a “risk level” that allows users to order results from most to least serious threats to their proposed mark. This image comes from Markify’s “comprehensive reports” while we used its “prosearch” for the comparisons between firms.

c) Trademarkia

Trademarkia is a visual trademark search engine that operates as a subsidiary of LegalForce, an intellectual property law firm.²³⁷ Trademarkia offers a number of trademark services including registration, legal action against infringing marks, trademark renewal, trademark revival, and trademark watch.²³⁸ Like other firms in this study, Trademarkia offers several services that can broadly be considered to be “brand management.”

Trademarkia distinguishes between “knockout” and “comprehensive” searches and offers both. Knockout searches comb the USPTO page for any similar marks, but these results do not guarantee that the identified marks are available or meet the standard for registrability.²³⁹ Its comprehensive search, on the other hand, furnishes users with a report that checks the mark against additional sources and contexts to ensure that the mark is available. It is a little unclear, but it seems that this process involves human input.

We focus on Trademarkia’s knockout searches, although we note that this service is free, and thus Trademarkia’s comprehensive search results may be better. However, the free service most closely resembles the other search engines in our study because it appears to use an algorithmic approach without human input.

237. See LEGALFORCE RAPC WORLDWIDE, <https://www.legalforcelaw.com/> (last visited July 28, 2019); see also *Trademarkia*, CRUNCHBASE, INC., <https://www.crunchbase.com/organization/trademarkia#section-overview> (last visited July 28, 2019).

238. See the “Trademark” drop down menu on <https://www.trademarkia.com/>.



239. For example, if one searched “google” as a trademark on Trademarkia’s free service, a note appears at the bottom of the search results that says:

NOTE: Trademarkia.com is updated regularly with the latest trademarks from the United States Patent & Trademark Office (USPTO). There may be marks that were removed from Trademarkia at mark owner’s request. Trademark search results are not indicative of the availability of the trademark. Applications requested through Trademarkia are evaluated by an attorney for the availability of the trademark. The Google trademark has a greater likelihood of registration if it satisfies the following conditions: (1) it is not confusingly similar to other marks, (2) it does not dilute a famous mark, (3) it is not generic or descriptive, and (4) if there are no unregistered, common law trademark holders that are using this trademark in commerce today.”

TRADEMARKIA, <https://www.trademarkia.com/trademarks-search.aspx?tn=google>.

Figure 6: Trademarkia's Search Results²⁴⁰

Click here to Register a Trademark »
Registration service fee is just \$199 + govt fee!

Mark Name	Description	Status
SERIES 1 SERIES 1 Filed on: 5/13/2010	binoculars; camera lenses, binocular lenses and telescope lenses Read More...	CONTINUED USE AND INCONTESTIBILITY ACCEPTED 7/7/2018 Live Mark
 SERIES 1 Filed on: 3/14/1991	loudspeakers Read More...	CONTINUED USE NOT FILED WITHIN GRACE PERIOD, UNREMOVABLE 3/22/1999 Dead Mark Register A Trademark
 SERIES 1 Filed on: 3/14/1991	stereo amplifiers Read More...	CONTINUED USE NOT FILED WITHIN GRACE PERIOD, UNREMOVABLE 12/29/2000 Dead Mark Register A Trademark

Chat with us!

d) TrademarkNow

TrademarkNow was founded in 2012²⁴¹ and is explicitly premised on using AI to revolutionize trademark search. It explains that its search engine,

[a]t its core is a unique artificial intelligence model of trademark law based on both explicit and intricate domain models of the law. Created by experts in trademark law and linguistics, our cutting-edge system also utilizes state-of-the-art machine-learning techniques to produce models that seamlessly take real-world complexities into account.²⁴²

Essentially, it tries to encode law, legal rules, and intuitions about legal interpretation of IP into its models in order to furnish users with the most relevant results.

TrademarkNow offers a few different products.²⁴³ ExaMatch is intended to be a first step for any trademark applicant and promises “instant screening”

240. Note that it provides descriptions and statuses in addition to the mark name.

241. *TrademarkNow*, CRUNCHBASE, INC., <https://www.crunchbase.com/organization/trademarknow> (last visited July 28, 2019).

242. *About Us*, TRADEMARKNOW, <https://www.trademarknow.com/about> (last visited July 28, 2019).

243. *See Brand Protection – NameWatch™*, TRADEMARKNOW, <https://www.trademarknow.com/products/namewatch> (last visited July 28, 2019).

results for trademark results.²⁴⁴ The company also offers a “clearance search” algorithm called NameCheck, which we used for our analysis, that improves upon knockout searches, and also a brand protection service called NameWatch that checks to see if anyone tries to register a conflicting mark.

3. *Our Methodology*

Although there are strong theoretical underpinnings for trademark search, there is little systematic evidence about how searches occur in practice. The fact that multiple products exist to assist potential registrants suggests that there is a real demand for tools that ease the trademark search process. In this section, we present results from a novel exploration of the efficacy of various trademark search engines. By comparing and contrasting particular results, we studied how well these search engines identify potential conflicts under Section 2(d) of the Trademark Act, 15 U.S.C. § 1052(d),²⁴⁵ which forbids the registration of a trademark that is confusingly similar²⁴⁶ to an existing registered trademark.

As discussed below, answering this broad research question turns on making choices about particular metrics. Our basic approach involved searching across each trademark search engine to evaluate how well each one picks up on potential conflicts. To address this question, we generated a list of “conflicted marks” that we knew should be flagged as a potential 2(d) violation. Second, using this list, we ran searches across each engine, and then measured the returned results. We then compared the results across several search engines using several relevant metrics.

Moreover, in the interest of reproducible research, we also created an end-to-end code pipeline. Each step of the process is entirely programmatic and can be easily reproduced by re-running the same scripts that we ran.²⁴⁷ The main advantage of taking this approach is that tasks like choosing conflicting marks involved no subjective judgment. Most importantly, automating searches allowed us to conduct this study at scale and collect data that would otherwise take an enormous amount of time and effort to record.²⁴⁸

244. *Preliminary Trademark Search – ExamMatch™*, TRADEMARKNOW, <https://www.trademarknow.com/products/examatch> (last visited July 28, 2019). TrademarkNow suggests that users, “[s]pend your time on the names that matter and not the ones that don’t.” *Id.*

245. 15 U.S.C. § 1052 (2018).

246. UP-COUNSEL, INC., *supra* note 210 (discussing “Likelihood of Confusion FAQ”).

247. Our code was mainly written using Python and R, and we will make it available upon request.

248. In this paper, we use about 100 different search terms. In previous work, Moerland and Freitas used terms related to just one mark, “Apple Inc.” *See generally* Moerland & Freitas,

To summarize, the basic methodology took the following steps:

- 1) We developed a list of conflicting proposed marks that should be flagged by a given search engine as being ‘confusingly similar’ to a preexisting mark.
- 2) Searched a term across all of the search engines and returned relevant results.
- 3) Saved all of the search results.
- 4) Repeated this procedure for each search term.
- 5) Analyzed the number of killer marks, precision/recall, and other metrics.

4. *Generating Conflicted Trademarks*

Generating a list of trademarks to run through search engines was a conceptually challenging task. We wanted to emulate the typical use case as much as possible when doing searches. Tackling this problem meant that the list of trademarks had to resemble actual searches that registrants would reasonably conduct.

One identified potential approach was to take a set of existing registered trademarks, either randomly chosen across goods or services or optimized for a particular trademark class,²⁴⁹ and search for them in each search engine. The attractive feature of this approach is that searched trademarks should almost definitely be flagged as problematic because there should be an exact match for them in TESS and other databases. If a search engine did not capture this conflict, it would be a signal of poor quality. Unfortunately, searching currently registered trademarks does not reflect how registrants actually use these trademark search engines prior to registration. Since registrants are looking to see whether their own mark is likely to run into a conflict, it is unlikely that anyone would search preexisting marks with any regularity.

Another identified approach is to create fake trademarks to search that closely resemble existing marks. For instance, one could swap a few letters in an existing mark to create a new mark, and then search the new mark to see if it would be flagged as confusingly similar to the original. Again, this approach seems attractive but does not reflect the true data generating process.

supra note 6. Our method, we think, provides a way to supplement qualitative studies like that of Moerland and Freitas.

249. See Brian Farkas, *Trademark Classes: Which One Fits the Mark You Are Registering For?*, NOLO, <https://www.nolo.com/legal-encyclopedia/trademark-classes.html> (last visited July 28, 2019).

Registrants likely create potential trademarks through creative processes and in ways that are associated with the brand they hope to protect. Swapping out letters and phrases would return marks that are confusingly similar to the originals, but not reflect how registrants actually create their own marks.

For these reasons, we instead scraped recent 2(d) rejections, generated a list of them, and used this list of marks for our searches. This approach overcame the fundamental flaw inherent in other approaches, namely that they do not reflect the actual creative process that generates confusingly similar marks. Additionally, by searching marks that were already rejected for being confusingly similar, we avoided needing to make personal judgments about what the USPTO might consider to be a 2(d) violation. For the same reason, by relying on a list of prior trademarks, we also did not have to occupy the minds of trademark registrants and try to emulate their thought processes when creating trademark names.

Ideally, we would have been able to observe the actual searches that registrants conduct across all search engines. In practice, however, generating this sort of list would be difficult because it would require each search engine firm to disclose its customers' identities, internal algorithms, and business practices in detail. We also had no information regarding whether the marks we searched—or their rejections—were in the datasets that the search engines had been trained on. In machine learning, separating a training set from a validation or test set is important because an algorithm can overfit to the training set, meaning it learns the patterns in that data but does not generalize well. Metrics like accuracy will seem artificially high if reported on the training set for this reason, and therefore a held-out validation/test is important for simulating how the algorithm performs when given *new* data. It is possible that the search engines in our study have seen the marks we search before, but it would be hard to quantify if this happened and whether the results would change.²⁵⁰ Using 2(d) rejections at least resembles the sort of marks we should expect a search engine to flag, and it avoids the pitfalls of trying to replicate the data generating process wholesale.

250. See Gareth James, Daniela Witten, Trevor Hastie & Robert Tibshirani, AN INTRODUCTION TO STATISTICAL LEARNING WITH APPLICATIONS IN R 176 (G. Casella, S. Fienberg & I. Olkin eds., 2017), available at <https://statlearning.com/ISLR%20Seventh%20Printing.pdf> [<https://web.archive.org/web/20210114184648/https://statlearning.com/ISLR%20Seventh%20Printing.pdf>] (explaining this reasoning).

5. *Scraping Websites*

Once we generated the list of conflicted trademarks, we turned to running them through each search engine. We wrote Python²⁵¹ scripts to achieve this task. Running searches programmatically has several advantages. The primary benefit is that the searches scale easily; conducting ten, a hundred, or a thousand searches requires no additional effort on the part of the analyst, simply more time to run queries. Future studies can therefore use this code as a template to expand upon, confirm, or adjust our results.

Another advantage is the ability to easily make multiple test runs to understand which configurations will get the best results. Search engines have several different search features such as searching for translations, including dead marks, or looking for different types of matches. Optimizing each search engine for its typical use case is key, and being able to run multiple tests easily helps with calibration.

Finally, creating reproducible scripts ensures transparency, which is critical when we are evaluating various software platforms. By enabling anyone to read the code, understand it, and replicate it to guarantee the accuracy of the results, we obviate concerns about mistakes in the research process. These concerns are further mitigated by drawing on the common tools Selenium and BeautifulSoup to complete our research.²⁵² Combining these two tools make it possible to create scripts that consistently and reliably scrape data from each firm in our study. With relatively simple code, it is possible to generate a rich dataset that allows us to answer a novel research question. Similar studies that

251. Python is a popular programming language in software engineering, data science, and other computer programming tasks. It is free to download and use. *See* PYTHON, <https://www.python.org/> (last visited July 28, 2019).

252. In terms of technical details, the primary tools we used were the Selenium and BeautifulSoup packages in Python. Selenium is a package that enables automated web browsing through a variety of common browsers. *See* SELENIUM, <https://selenium.dev/> (last visited July 28, 2019). Using Selenium, it is possible to automatically navigate to a trademark search website, login, and run search terms. The basic principle underlying Selenium is that if it is possible for a human to click or enter text in any part of a website, it is possible to automate this process with Selenium. The major drawback of Selenium is that if a website's underlying source code changes, then it could potentially break a webcrawler. BeautifulSoup is another common package that can take a webpage and break down its HTML in a convenient format for humans to read. *See* *Beautiful Soup*, CRUMMY, <https://www.crummy.com/software/BeautifulSoup/> (last visited July 28, 2019). The main feature here is that it provides HTML tags for every element on a webpage. In our case, this feature makes it easy to scrape tabular or list results for each of our search terms in an automated fashion.

look at other areas of law (whether in IP or otherwise) could easily replicate our general approach.

The pseudocode for accomplishing these steps is as shown in Figure 7.

Figure 7: Pseudocode

This program navigates to a trademark search engine, loops through a list of trademarks, searches each mark, and returns a table in a dataframe.

1. Load the list of “conflicted trademarks” to be searched
2. Initialize an empty dataframe object
3. For each trademark in the conflicted trademarks list:
 - a. Try:
 - i. Initialize a Selenium webdriver
 - ii. Navigate to search engine website
 - iii. Find the “search bar”
 - iv. Enter trademark
 - v. Extract XML page source
 - vi. Extract table
 - vii. Save table to a dataframe
 - viii. Click to next page and repeat steps i–vii if necessary to build table sequentially
 - b. Except:
 - i. Pass

6. *Exploratory Data Analysis*

At a basic level, we are interested in whether a search engine provides the user with adequate information to dissuade them from attempting to register a mark that is likely to get 2(d) trademark rejection. However, it is not straightforward to pick a single metric that satisfies this proposition. For this reason, defining metrics is a key task because the core research question could be interpreted in many different ways.

Before turning to an evaluation of the engines, we first provide some exploratory data analysis to build intuition around trademarks, search engines, and the notion of “similarity.” For purposes of our study, we identify two major categories of conflicts: either confusingly similar spelling or sound. Since it is easier for a search engine or other algorithm to detect spelling similarity, rather than phonetic similarity, we expect that all search engines will have lower scores on the latter. Still, potential registrants are likely to be concerned with both types of conflicts, thus making the breakdown useful for comparisons between the search engines.

In our exploratory data analysis, we recorded the number of exact matches, as well as the number of phonetic matches. For exact matches, we checked whether the result is exactly the same as the searched mark. For phonetic matches, we employed the Soundex algorithm. The Soundex algorithm matches sounds by taking the first letter of a word, and then encoding the remaining consonants according to a predefined schema that assigns particular letters to particular number values.²⁵³ Note that Soundex is a fairly simple algorithm that is prone to errors, and the search engine's own algorithms are undoubtedly more sophisticated. That being said, we include it mainly to illustrate the concept of phonetic similarity.

We also looked at the total number of results returned by a search engine. Typically, a lawyer is obliged to look through each search result before making a recommendation to a client. Too many results can create unnecessary noise and add to the search costs, but too few results can put the user in the position of mistakenly filing a bad mark. By itself, the number of results is not too informative, but may be useful when put into context with other firms' results. We also provide breakdowns for results tagged as "high" or "low" risk (or the equivalent, whenever available).

a) Baseline

For our baseline, we use the USPTO's TESS results. This choice is a natural baseline because searching TESS is a typical first step for most applicants; it is freely available, and the database is directly connected to the USPTO's own information that it uses in decisions to approve or deny a trademark. By treating the TESS results as a baseline, we implicitly assume that paid services should do better on at least some measures.

The basic results pulled from TESS are in Table 1.

253. For more details, see *Soundex System*, NATIONAL ARCHIVES, <https://www.archives.gov/research/census/soundex.html> (last updated May 30, 2007).

Table 1: TESS Results Example²⁵⁴

marks	conflicted_mark	exact_match	levenshtein_distance
EX SERIES ZONE	SERIES 1	FALSE	8
TOUGHY U-GATE SERIES	SERIES 1	FALSE	15
ECONO GATE SERIES	SERIES 1	FALSE	12
ATLANTIS G-GUTTER CANOPY SERIES	SERIES 1	FALSE	26
KITCHEN ESSENTIALS CLUB SERIES	SERIES 1	FALSE	25
EVERY PRIZE ONLY A. AUCTION L THE UNITED	SERIES 1	FALSE	254
METAS -SERIES	SERIES 1	FALSE	9
LEGACY SERIES	SERIES 1	FALSE	9
CHAMPIONS OF EDUCATION MVP SERIES	SERIES 1	FALSE	30
MVP SERIES CHAMPIONS OF EDUCATION	SERIES 1	FALSE	27
FEDERAL RESERVE BLUNT JMJ THE UNITED	SERIES 1	FALSE	265
SQUATTERS SECRET STASH -OFF SERIES	SERIES 1	FALSE	28
SYMETRA LINK FIXED INDEX ANNUITY - SERIES	SERIES 1	FALSE	36
DR LIONIS: SO NATURAL BUSINESS FIRST AND	SERIES 1	FALSE	2070
SERIES	SERIES 1	FALSE	3

Other search engines might provide even more information than what we study in our paper. For example, whether a mark is live or dead, high/medium/low risk for a violation, and owner information might be provided as well. While this information could be interesting to explore, it is not necessary to answer the core question of how well trademark search engines flag potential 2(d) violations.


Our basic measure for efficacy is the “exact match.” An exact match corresponds to an instance where we search a mark that we know was rejected under 2(d), and then see if that exact mark is already registered. In the above example, we looked to see whether the mark “SERIES 1” already exists. An exact match is the most straightforward and least subjective way that a mark can conflict with a preexisting one. The one caveat to this statement is that two marks may share a name if they belong to entirely separate classes, and therefore are unlikely to degrade the quality of the other. In such instances, an exact match does not necessarily result in a rejection.

In terms of exact matches, we show results in Figure 8. In this sample, it is clear that it is fairly uncommon to recover an exact match. About a quarter

254. Here, we used the term “marks” to correspond to a returned result; “conflicted_mark” to refer to a mark that was previously rejected under 2(d); “exact_match” denotes whether, for a given row, the value in “marks” corresponds to the value in “conflicted_mark.” “Levenshtein_distance” refers to the edit distance between “marks” and “conflicted_mark.” We calculated exact_match and levenshtein_distance columns ourselves. The basic procedure is that we would take each of the “conflicted_mark” values (like SERIES 1), search them through each search engine, then store all of the search results as “marks,” and calculate these measures.

of the results are exact matches. What makes this figure interesting is that it provides some evidence that trademark search is a more complicated process than simply looking for whether one's proposed mark already exists. Rather, most potential conflicts will not match exactly and therefore require some judgment about likelihood of confusion.


Figure 8: Exact Matches in TESS



Our next measure was for “phonetic matches,” as shown in Figure 9. Using the Soundex algorithm,²⁵⁵ we looked for whether two strings match based on a phonetic encoding. These results provided more matches, with about half of the results returning a positive phonetic match. Again, the interesting bit here lies in the results that were not matches; discerning whether there is a signal in this group of marks is a key task for any improvements over TESS.

255. *Soundex System*, *supra* note 253.

Figure 9: Phonetic Matches in TESS



Finally, in Figure 10, we look at the overall number of results returned for each searched mark. This is an important metric because it contains a few key pieces of information. A large number of results could imply that a search engine did a good job exhausting all possible conflicts and returning a lot of relevant information. On the other hand, a large number of results could also imply that a search engine produced a lot of noise, perhaps too much for a human to reasonably sift through. Below, we visualize a random sample of searched marks and the number of results returned for each mark. Again, this is a random sample so one should not draw an inference from the shape of the distribution. However, it does provide a useful baseline for what a registrant can expect to find when they search TESS.

b) Exploratory Analysis of Private Search Engines


Below, we present results that compare how AI-powered search engines fare compared to the USPTO's own TESS system. Like above, our basic metrics are number of search results returned, how many exact matches are returned, how many phonetic matches are returned, and how many "close" matches are returned by letter substitution.

The overall takeaway from the results is that AI-powered trademark search engines indeed provide valuable insights for potential applicants. Either by pulling in additional information or packaging it in more manageable ways, they, in general, improve over the baseline results in interesting ways. It is also worth noting that they optimize for different things and thus may be better suited to different use cases.

Consider, for example, the number of matches that each search engine returns. Figure 11 illustrates this point. We searched 115 different marks, and Markify returned around 27,000 potential matches across these, while TESS and Trademarkia returned about 8,000, and TrademarkNow returned approximately 3,000. This comports with expectations, since Trademarkia and TrademarkNow seem to heavily rely on cross-checking against TESS, while Markify pulls in additional sources.²⁵⁶

256. MARKIFY, *About Markify*, *supra* note 233.

Figure 11: Number of Matches by Search Engine



Digging deeper, we can also see this difference visualized across different search terms. Figures 12–16 illustrate the number of search results per mark for each private search engine in our study. Note for these figures, we sampled ten marks to visualize the data. In general, each engine that returned a similar number of relevant results, often between 10 and 20. Certain marks, however, returned a much larger number of results.

Figure 12: Search Results by Search Term in TESS




Figure 13: Search Results by Search Term in Markify




Figure 14: Search Results by Search Term in Trademarkia




Figure 15: Search Results by Search Term in TrademarkNow






Figure 16: Search Results by Search Term in Corsearch



Investigating further, we can see the utility of AI-driven search by looking at our simplest metric, exact matches. TESS, Corsearch, Trademarkia, and TrademarkNow all returned a similar number of exact matches across all searches. However, if a trademark applicant was optimizing solely on finding exact matches, they might prefer a private search engine. Note that in Figure 17, TESS returned many more results that are not exact matches, while Trademarkia and, especially, TrademarkNow filtered out much of this noise. The AI systems underlying both of these private search engines aim to return fewer results overall, and thus better amplify the signal provided by the actual “exact matches” in the data.

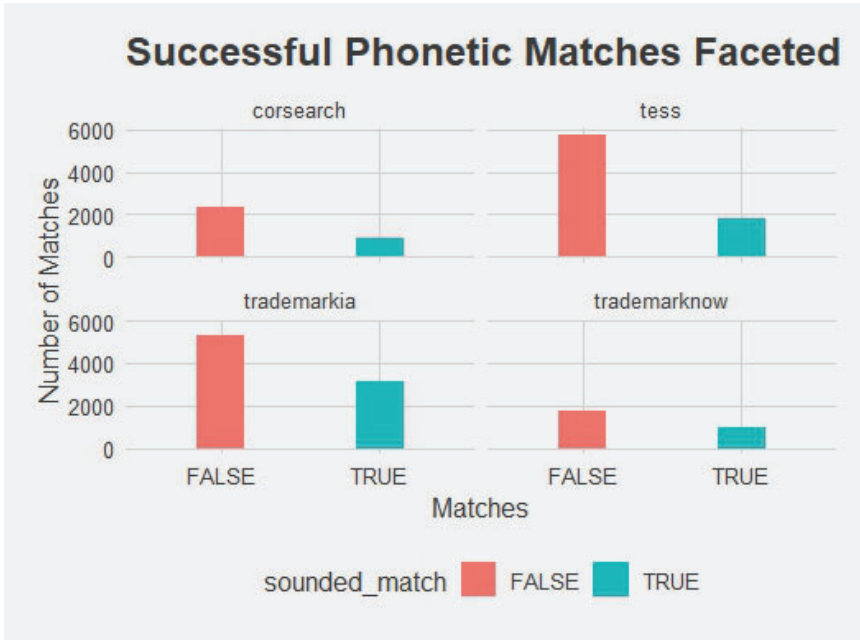
Figure 17: Exact Matches in TESS, Trademarkia, and TrademarkNow



We see a similar phenomenon with phonetic matches as well. Note that in Figure 18, each search engine tends to recover more phonetic matches than exact matches. This result is unsurprising as phonetic matching is less restrictive than exact matching. Also note that Trademarkia does remarkably well as nearly 40% of its results actually match phonetically. Similarly, TrademarkNow achieves nearly 50%. Relative to the TESS baseline, these AI-driven results represent a substantial reduction in noise.

Interestingly, Corsearch, which specializes in phonetic searches, seems to achieve a great deal of noise reduction. It returns fewer results overall than TESS, Trademarkia, and TrademarkNow, but generally returns a higher proportion of phonetic matches. This implies that the algorithm filters out a lot of irrelevant results and does a fairly good job prioritizing actual phonetic matches. Again, Corsearch’s own phonetic match algorithm may differ from the Soundex algorithm’s rules, so our results may understate the extent to which it successfully finds phonetic matches.

Figure 18: Successful Phonetic Matches Across Search Engines



The major takeaway from these results is that AI truly is transforming the trademark search landscape. Even on these basic metrics of exact and phonetic matches, a trademark applicant has little reason to use TESS over a private competitor, particularly when some of these private search engines offer their basic search functions for free (and charge for brand management instead), offering substantial efficiency gains. By returning fewer results in general and successfully filtering out irrelevant results, they make it easier to find knockout conflicts. Basically, this algorithmic approach achieves significant noise reduction at virtually no additional cost to the user.

7. Metrics

For our main results, we focused on whether a search engine successfully finds the mark that the USPTO cited in its 2(d) rejection (i.e., the “killer mark”). A killer mark is essentially an existing trademark that justifies rejecting a new application. If a search engine successfully uncovers such a mark, then it succeeded in providing the applicant with information about whether their proposed mark will be accepted. If the search engine fails to find this killer mark, the probability that an applicant goes ahead with a frivolous application rises.

To examine whether the search engines in our study successfully find the killer marks, we used the following metrics. For any given search result, we checked:

True Positive: The search result matched a killer mark

False Positive: The search result did not correspond to a killer mark

False Negative: There was a killer mark that did not have a match in the search results.

Conceptually, these metrics are usually presented alongside “True Negatives.” However, we cannot identify true negatives in this context because that would correspond to no search results returned and no killer marks present. That being said, we can still further combine the preceding metrics in useful ways:

Recall: Ratio of killer marks found to total killer marks, i.e., True Positive / (True Positive + False Negative)

Precision: Ratio of search results that were actually killer marks, i.e., True Positive / (True Positives + False Positives).

These metrics are frequently used in machine learning for classification problems and work well in this context, too, because they can give us a sense of how each search engine performs, and the tradeoffs among them. For instance, one search engine may prioritize recall (i.e., finding all of the relevant killer marks) over precision (i.e., not flagging false positives), or vice versa.

In the results section, we present these metrics in a few different ways. First, we tweak these definitions slightly to see how well each search engine does at finding any killer mark (instead of all of the killer marks). We then show precision and recall for all search results in our overall dataset. Finally, we show the same metrics for when we limit the number of returned search results per trademark application. Note that in calculating these numbers, we only used results from each search engine’s basic search that was the equivalent of a “knockout” search.

8. *Results*


Our results suggest that the landscape of trademark search is rich and interesting, and there is a real potential to further study search costs borne by trademark registrants. The main takeaway is that private trademark search engines provide a genuine value-add to a potential trademark registrant. While not all private search engines provide a meaningful improvement over free,

public options, there is evidence of meaningful differentiation between various products.

Our exploratory analysis illustrates some of the basic questions in trademark search. Specifically, we showed that there is some evidence of differentiation between different search engines and the USPTO’s own search engine. Differences in number of results returned, the types of matches, and other features may be relevant. In this part, we look at how each search engine performs with respect to our precision and recall metrics to examine these differences in greater depth.

Before delving directly into precision and recall, we first look at whether a search engine finds at least one killer mark associated with a particular trademark search. Figure 19 shows the number of instances in which a search engine finds at least one killer mark. In this case, TESS actually does not perform so poorly relative to private sector search engines. In general, most search engines fail to find a killer mark more often than not.

Figure 19: Killer Marks Found by Search Engine



However, we calculate precision and recall somewhat differently. Instead of asking whether a search engine finds at least one killer mark, we instead ask,

“Of all of the killer marks in the dataset, how many was each search engine able to detect?” Some searched trademarks have multiple killer marks associated with them, so precision and recall here will capture whether a search engine uncovered all of the relevant killer marks.

In Table 2, we show results where we derive the precision and recall for each search engine, without limiting the number of results that each search engine returns. Results show that every private search engine achieves higher recall than TESS, and many improve on precision as well.

Table 2: Precision-Recall of Search Engines

Search Engine	Recall	Precision
Corsearch	0.369919	0.028086
Markify	0.609756	0.006916
Tess	0.146341	0.006524
Trademarkia	0.105691	0.04878
TrademarkNow	0.691057	0.06308

We can also explore how a potential applicant could tradeoff between recall and precision. Figure 20 shows the information from Table 2 as a scatter plot with recall on the y-axis, and precision on the x-axis. Figure 21 shows the same information, but this time with varying limits on the number and results per search for each search engine.

Figure 20: Precision and Recall


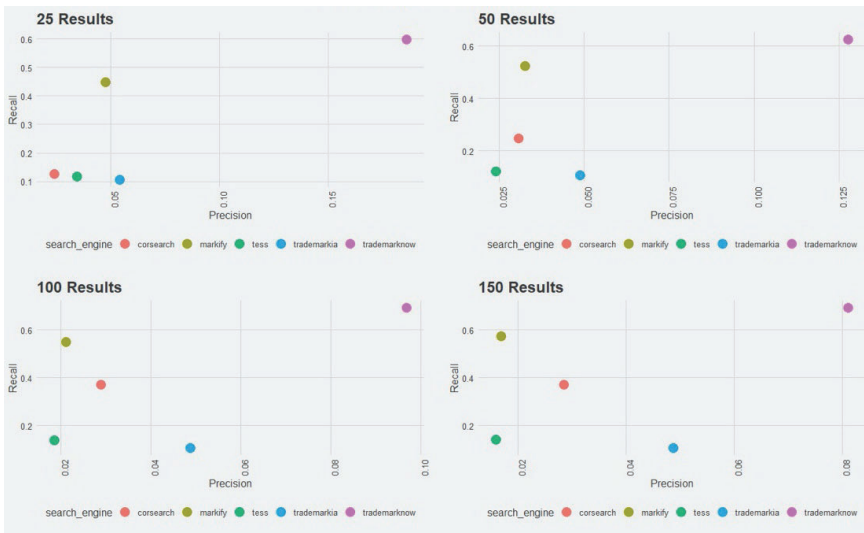


Figure 21: Precision and Recall with Limited Search Results



Interestingly, different search engines exhibit different behaviors depending on the limit on the number of search results. Corsearch's precision tends to improve with additional results, while its recall becomes stable around 0.40. Trademarkia also tends to improve on both measures with additional results, and tops out on recall around 0.20. Markify consistently achieves recall in the 0.55–0.60 range, but loses precision with additional results. TrademarkNow similarly achieves a similar recall across the board, but maintains a higher precision than any of the other search engines, even though it decreases with additional searches. More simply, one could stop searching after fifty results and likely already have found the killer mark in Markify and TrademarkNow, while that threshold may be closer to a hundred in Corsearch and Trademarkia. These numbers should also be properly understood as an estimate taken from a sample in a particular time. The marks were scraped in 2019 and searched between 2019 and 2020. Since these are dynamic systems, differences in the sample or time could change these findings considerably.

Taken together, these results indicate that the search engines prioritize certain marks in their search results. Some like Markify and TrademarkNow make this explicit with riskiness indicators.²⁵⁷ Others seem to do such ordering more implicitly. If precision and recall both increase with additional searches, that indicates that killer marks tend to be identified somewhere other than the beginning of a search result list. On the other hand, if recall remains stable and precision decreases, that indicates that the search engine already found the relevant killer mark.

Both the number of times that a search engine finds a killer mark and the precision and recall scores yield valuable insights. Even one killer mark would be enough to defeat a trademark application, so successfully finding at least one is important for assisting potential trademark registrants. Most trademark search engines do about the same as the USPTO on this measure, and all search engines perform better than the USPTO at finding all possible conflicts.

IV. IMPLICATIONS FOR FURTHER STUDY

Today, AI is rapidly reinventing the process of search altogether, particularly in areas of law and government. Court cases, congressional hearings, and government documents are all examples of areas where AI may soon be used in search tools.²⁵⁸ Already, AI is being deployed to search

257. Markify uses “high risk” and “low risk” classifiers, while TrademarkNow shows the percentage likelihood of riskiness.

258. Faraz Dadgostari, Mauricio Guim, Peter A. Beling, Michael A. Livermore & Daniel N. Rockmore, *Modeling Law Search as Prediction*, ARTIFICIAL INTELL. L. (2020), <https://>

databases of parking tickets for those who want to contest them.²⁵⁹ Within the world of IP, we see AI-related techniques throughout the global marketplace, and more and more countries and companies have turned to the tools of machine learning to refine their techniques.

These AI-powered techniques are especially important, not just for the purposes of refining search, but also because of the insights they offer into the economics of IP. On a scholarly level, as our comparison shows, a new area emerges for future research on firm search costs within the trademark registration system through the intersection of AI and trademark search processes. In this Article, we showed how AI is revolutionizing the economics of search in the trademark space, raising new questions about the role of AI in brand management more generally. The main implication of our research is that search costs and AI will continue to be important to legal decisions, both within IP and outside of it. As we showed, firm search costs are a dramatically overlooked area of study and may ultimately hold the key to studying the role of AI in trademark law.

A. OUTCOMES AND IMPLICATIONS

As we have suggested, when scholars and practitioners explore the potential role of AI in transforming patent prosecution and litigation, they may also benefit from looking at trademarks. Trademarks are incredibly valuable assets, and studying their role in the AI-powered marketplace reveals core insights into the economics of IP system at large. As we have shown, AI carries the ability to efficiently compare a proposed trademark against millions of registered trademarks and to assist in determinations about the proposed trademark's worthiness of protection. As with patent and copyright infringement, effective deployment of AI tools prior to the creation of a property right in the trademark context could substantially reduce litigation and other costs when real conflicts arise later on.

At the outset, our legal system places the core responsibility for trademark search and enforcement on the trademark holder.²⁶⁰ Thus, one interesting question that may be worth exploring is how the use of AI in the USPTO context compares to other government contexts. As we have noted throughout this piece, trademark search engines largely emerged because the USPTO does not enforce existing trademarks against potential conflicts.

doi.org/10.1007/s10506-020-09261-5 (suggesting a model of law search based on a notion of search space and search strategies).

259. Shannon Liao, "World's First Robot Lawyer" Now Available in All 50 States, THE VERGE, (July 12, 2017), <https://www.theverge.com/2017/7/12/15960080/chatbot-ai-legal-donotpay-us-uk>.

260. USPTO, *supra* note 188 ("Trademark Basics").

Trademark owners are responsible for discovering and taking legal action against potentially damaging marks.²⁶¹

As such, optimizing the search process benefits both the trademark holder as well as the overall marketplace for trademarks generally, ultimately benefiting consumers. If the USPTO grants too many confusing trademarks, then the market would produce weaker trademarks, harming consumers in the marketplace and leaving more marks vulnerable to enforcement by others. From a registrant's perspective, avoiding a potential rejection saves a lot of time and effort that would otherwise be wasted, conserving the strength of the mark that is ultimately registered. As we have argued, the crucial moment of initial search is a key part of the brand-creation and management process, forming an important threshold of protection.

In this study, we looked at one aspect of the trademarking process—the search for possible conflicts prior to registration—and the significance of search in terms of rethinking our approach to trademark law altogether. The trademark search technologies that we studied here are some examples of how new computational techniques are attempting to solve this puzzle by modeling human decision-making. To summarize, AI lowers search costs by doing a lot of the hard work of making substantive inferences about the relationships between different trademarks, thus empowering applicants to make informed decisions about whether to proceed with their trademark applications. Engstrom et. al. provide an in-depth look at the USPTO's current experiments with AI adjudication, specifically in the realm of patent examination.²⁶² They note that AI has the potential to reduce search costs for the examiners, but thus far has not been fully implemented as the tools mostly improved the work of examiners with computer science backgrounds.²⁶³ Our study suggests that the development of private sector alternatives in the trademark space might make these tools more broadly accessible. Indeed, the UPSTO is currently exploring implementing deep learning models on image searches.²⁶⁴

Ultimately, as we suggest below, our exploration of trademark search engines and the choices we made with regards to methodology and metrics could have interesting lessons for other similar studies in different areas of law. Our study also revealed some important conclusions about the process of trademark registration and the important role that search costs can play in the process.

261. USPTO, *supra* note 164.

262. *See generally* Engstrom et al., *supra* note 64.

263. *See generally id.*

264. *See generally id.*

The first main takeaway from this exploration is that AI is already being used in this space, and it is capable of reducing search costs through algorithmically driven information retrieval. Noise reduction and algorithmic prioritization are two major features that these AI search engines achieve. Trademark applicants now have access to tools that can process millions of preexisting trademarks, analyze them, and produce relevant outputs that human beings can understand.

Second, consistent with the literature that finds that consumers give significant weight to non-monetary attributes (like brands, reputation, service quality and pricing quality) in making purchase decisions,²⁶⁵ we found that trademark registrants, in using AI-powered search, also enlist a variety of non-monetary variables in their own considerations, such as trademark class and lexical similarity to existing marks. This means that search engines can optimize on many more variables than just trademark strength alone. At the same time, it is reasonable to presume that the addition of non-monetary elements, such as the ones that we have seen, can play a determinative role in the trademark registrant's selection of a search engine. Some of these non-monetary attributes may turn on the risk of litigation, the magnetism of the mark, or the mark's relationship to other identities and marks, among others.

A third takeaway involves optimizing the prediction of the outcomes of both registration and potentially litigation. Our results provide some basic validation of the central premise that these types of legal outcomes can be mathematically modeled. These models can detect lexically and phonetically similar marks and, importantly, can sift out results that do not meet certain similarity thresholds. Some attach explicit risk scores, while others implicitly calculate them and then order results. As expected, these risk determinations may follow expected patterns both in distributional shapes and over time, but the patterns may change in the future as adversarial models develop.

None of the trademark search engines we studied model whether a mark objectively meets or fails to meet the 2(d) standard. Such an objective truth plainly does not exist. Rather, these search engines attempt to model the ways that the trademark office, or rather the people in the trademark office, reach their determinations. Implicitly, by making choices about which marks to return and ordering them in a specific way, these search engines make the claim that they can approximate trademark examiners' decision-making well enough to guide trademark applicants' and registrants' business decisions.

Finally, it bears mentioning that including a selection of a larger number of AI-driven variables in a trademark selection decision also introduces the

265. See Zhang et al., *supra* note 23, at 91.

potential for an extremely complex decision-making process. Conceptualizing the potentially unlimited set of variables is practically impossible. The scale of the trademark search space is also massive with millions of registered trademarks in existence. Traversing this massive set of trademarks and retrieving the ones that could present potential conflicts implies enormous search costs for registrants and trademark examiners alike. Search engines try to ameliorate these costs by reducing noise. But, as we have shown, some search engines are better than others at reducing the level of noise encountered by applicants and correcting for the information asymmetries that arise. Specifically, these search engines optimize on certain similarity metrics and drive their results with them. Today, it is difficult to surmise how these AI-driven effects might play out in trademark litigation, i.e., whether they would increase or decrease the occurrence of litigation or its costs. On this point, more research will be needed in the future.

B. FRAMING TRADEMARK REGISTRATION AS AN ADVERSARIAL MACHINE LEARNING PROBLEM

As these AI search tools mature, we expect that trademark registration will start to resemble an “adversarial machine learning” problem.²⁶⁶ Previous literature in IP and administrative law identified the back-and-forth between the USPTO and patent applicants.²⁶⁷ These pieces discussed the problem of the PTO adapting to increasing sophistication in patent applications.²⁶⁸ This sophistication is in part driven by the use of AI tools, and, in turn, the USPTO might consider using machine learning to improve its own capacity to conduct meaningful examinations.²⁶⁹ Because applicants have strong incentives to maximize the scope of their claims and the USPTO has an incentive to minimize this scope, the two sides will each construct their decisions in anticipation of the other’s incentives.²⁷⁰

To build on this literature, we suggest also reframing trademark search as an adversarial machine learning problem. Adversarial machine learning refers to machine learning applications where underlying data distributions change in response to external stimuli. For instance, one problem in training AI for self-

266. For background on adversarial machine learning, see generally Ling Huang, Anthony D. Joseph, Blaine Nelson, Benjamin I.P. Rubinstein & J. D. Tygar, *Adversarial Machine Learning*, AISEC ’11 (Oct. 2011), <https://dl.acm.org/doi/pdf/10.1145/2046684.2046692>.

267. See generally Rai, *supra* note 3; Ebrahim et al., *supra* note 3, at 1193–95 (describing the inventor-examiner interaction).

268. See generally Rai, *supra* note 3; Ebrahim, *supra* note 3, at 1195–1211 (discussing the automation applications in patent prosecution).

269. See generally Huang et al., *supra* note 264.

270. See generally Ebrahim, *supra* note 3.

driving vehicles is that these AI systems can be easily tricked with just a little additional noise.²⁷¹ A self-driving vehicle may be trained to recognize a stop sign with high accuracy, but may suddenly fail if a stop sign has a sticker on it. Although a human being would still recognize the stop sign as such, the AI can be easily fooled because it has never seen this sort of example before.

To address this problem, an analyst may try to present the AI with “adversarial” examples in the training phase so that it can learn from these examples. In the self-driving vehicle example, this process could involve perturbing pixels in an image or providing examples of stop signs with stickers and other idiosyncratic markings. Thus, the AI can learn to improve its predictions, even when there is noise present.

Extending this concept into the trademark space, we can conceptualize the general problem articulated by authors such as Rai and Ebrahim in these terms. Consider the following theoretical model: Assume there was a universe of trademark applications prior to the advent of private trademark search engines. Once AI trademark search tools were built based on historical PTO decision data, the recommendations produced by these tools likely influence the names and types of marks in applications to the PTO, thus changing the underlying distribution of trademark applications.²⁷² The PTO, in response to this change, adjusts its own algorithms and procedures. The search engines retrain their models based on new PTO decisions, and, once again, influence the sorts of trademark applications that are eventually filed. And the PTO again must update its decision-making. This interplay between the PTO and trademark search engines (and trademark applicants) thus evolves dynamically over time.

By framing trademark registration as an adversarial machine learning problem, it becomes clear that the introduction of AI into the process of trademark registration also changes the substance of trademarks. When the PTO makes a series of decisions that search engines must retrain their models on, this represents the PTO adding new noise into their systems. Similarly, when trademark applicants file new applications that are optimized by advice provided by search engines, they add noise to the PTO’s decision-making. This dynamic game implies that, over time, the substance of applied for and registered trademarks may keep changing.

271. Solving the problem of malicious signage in particular is an active area of research in computer science. *See generally* Chawin Sitawarin, Arjun Nitin Bhagoji, Arsalan Mosenia, Mung Chiang & Prateek Mittal, *DARTS: Deceiving Autonomous Cars with Toxic Signs*, ARXIV.ORG (May 31, 2018), <https://arxiv.org/pdf/1802.06430.pdf>.

272. Indeed, we mention these selection effects as a hurdle for studying the causal effect of trademark searches in our methodology section.

Previous scholars have advocated for the use of machine learning in patent examinations as a response to increasing sophistication in the private sector. Adversarial machine learning makes clear why this call is important. The term “adversarial” may imply that the contest between the PTO and private sector search firms is damaging, but it should instead be thought of as a framework that improves the quality of trademarks and administrative decision-making. Here, the outputs of the PTO’s decisions become the inputs of the search engines’ algorithms, and vice versa. By dynamically responding to each other, the substance of trademark applications will change over time, and, ideally, in a way that gradually eliminates “easy” cases. Moerland and Freitas argue that so far, government search engines have not developed a level of sophistication that can replace human examiners. Future work might explore whether this argument holds true for such “easy” cases, or whether it is more applicable to “hard” cases involving novel or ambiguous marks.

Using adversarial machine learning as a model, we can open up new areas of inquiry in addressing situations where a public agency needs to make decisions based on information provided by a private actor. Adversarial machine learning provides a framework for thinking of dynamic government decision-making systems as responding to added noise. Just like adding random pixels to an image stress tests the AI system that powers a self-driving vehicle, policymakers can think about ways to utilize stress tests provided by private actors to better calibrate law, policy, and administrative decision-making. Thus, administrative agencies investing in machine learning tools and, more importantly, adopting theoretical frameworks about dynamic decision-making can empower them to improve over time.

C. RISK ASSESSMENT IN THE TRADEMARK ECOSYSTEM

Our study suggests that a supply-side study of trademarks should engage further with the search costs associated with post-registration enforcement, as well as the search costs inherent in the entire brand management process. Getting a trademark registered is important, but the post-registration landscape of enforcement is perhaps even more important. The largest question, perhaps for a future round of research, concerns the impact of AI on the overall trademark litigation ecosystem, i.e., whether or not search costs may have a similar effect on the trademark system like the patent system, where patent trolling and patent pooling have detrimentally affected the marketplace of patent acquisition and enforcement. With millions of existing trademarks spread across a variety of industries, it is simply infeasible to manually look for potential conflicts and deal with them as they arise. Instead, AI-powered tools can consume this tremendous amount of brand-related data, process it, and present it to the brand owner in a way that filters out noise while giving

trademark owners a way forward. In sum, by substantially reducing the costs associated with search, these tools also bolster trademark holders' abilities to protect their IP effectively.

One central question that can be raised from this project is similar to questions raised regarding the use of AI in other contexts: will AI transform trademark law altogether? Of course, given the rapid increase in trademarking activity in the past few decades, one can certainly understand the intuitive appeal of employing a greater use of AI. However, as Gangjee notes, “[t]he seductive appeal of the all-seeing algorithm should be resisted,” because it faces, at best, a current set of limitations.²⁷³ We believe, like other AI trademark experts, that while AI has the capacity to refine and improve the process of trademark search and registration, at its best, it should serve to complement, rather than replace, human judgment.²⁷⁴ Of course, it would be unrealistic to predict that AI-driven judgment will somehow diverge widely from human judgment, mainly because AI is normally trained on decision-making data that is generated by humans. As Gangjee notes, “where the data for a machine learning approach is derived from judicial content analysis—past decisions by human tribunals where factors are coded and correlations derived—the algorithm will behave like the human decision maker it is modelled after, warts and all.”²⁷⁵

In sum, as our paper has suggested, searching for preexisting trademarks is simply the first step in the process of overall brand management. While we conducted an in-depth look at the search process inherent in the trademark process, the platforms we studied also provide brand management services. Such services provide us with a deeper set of variables that may even go beyond the systems of patent pooling and enforcement that we have seen thus far. The same AI and machine learning tools that power their search engines also power their brand management tools, suggesting that further study of brand management and AI may be warranted.

Consider, for example, the rich set of possibilities that stem from providing a preliminary analysis of risk assessments in trademark search. Much of the existing literature that explore the use of AI-driven risk assessments in government decisions focus mainly on actors with enforcement powers in either criminal justice or administrative law. So far in legal, computer science,

273. See Gangjee, *supra* note 6, at 15.

274. *Id.* at 11 (adopting this view and quoting COMPUMARK WHITE PAPER, ARTIFICIAL INTELLIGENCE, HUMAN EXPERTISE: HOW TECHNOLOGY AND TRADEMARK EXPERTS WORK TOGETHER TO MEET TODAY'S IP CHALLENGES 5 (2018) (observing that AI is “intended to complement, not replace, human analysts”).

275. See Gangjee, *supra* note 6, at 11.

and policy literatures, discussions on the use of risk assessment in public policy primarily focus on the implications of AI tools on values like fairness, accountability, and transparency.²⁷⁶ Risk assessment has been the subject of debate in criminal justice, especially, with applications to sentencing,²⁷⁷ parole decisions,²⁷⁸ and bail reform.²⁷⁹ Scholars have also recently focused attention on critical issues like housing and employment, thus extending discussions on fairness in machine learning to include anti-discrimination and equal protection law.²⁸⁰ These discussions largely center around the legal problems and implications stemming from the use of “black-box” algorithms in decisions.²⁸¹ In particular, the scholarly community is deeply engaged with the possibility that algorithms can learn and reinforce human biases in a way that creates inequitable outcomes for marginalized communities.

Our results suggest that a ripe area for future research could be the use of risk assessments in IP law. Arti Rai describes the theoretical potential for the use of machine learning models in patent applications, and critically notes that many of the equity and justice concerns inherent in areas like crime and housing may not apply to IP contexts in the same way.²⁸² Given that the stakes are quite different, IP may be a good subject to explore and experiment with risk assessments in legal decision-making. This is especially because the government does not bear the same set of enforcement responsibilities in trademark law.

Engstrom et. al. have explored the idea of surveying the use of AI across government administration.²⁸³ They created a typology of different AI use cases in government such as enforcement, regulatory research, and adjudication.²⁸⁴ Through the exercise, they defined adjudication specifically as, “[t]asks that support formal or informal agency adjudication of benefits or rights,” and note that patent and trademark office applications as an

276. See Solon Barocas, Moritz Hardt & Arvind Narayanan, FAIRNESS AND MACHINE LEARNING, <https://fairmlbook.org/> (last updated Dec. 6, 2019, 3:49 PM).

277. See John Monahan & Jennifer L. Skeem, *Risk Assessment in Criminal Sentencing*, 12 ANN. REV. OF CLINICAL PSYCHOL. 489 (2016).

278. See Megan Stevenson, *Assessing Risk Assessment in Action*, 103 MINN. L. REV. 303, 304 (2019).

279. See generally Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig & Sendhil Mullainathan, *Human Decisions and Machine Predictions*, 133 Q.J. ECON. 237 (2018).

280. See generally Solon Barocas & Andrew D. Selbst, *Big Data’s Disparate Impact*, 104 CALIF. L. REV. 671 (2016).

281. See generally Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 Fordham L. Rev. 1085 (2018).

282. See Rai, *supra* note 3.

283. See generally Engstrom et al., *supra* note 64.

284. See generally *id.*

example.²⁸⁵ Our study suggests a general approach and method for assessing the interplay between the government's adjudication system and the private sector, and this general framework could also be applied to other areas of law as well. Zoning, licensure, and social security benefits claims are all examples of the government adjudicating the benefits and rights of private parties, and the ability to assess how AI-driven systems work in these spaces will likely be a rich, new research area.

With respect to trademark law, our work suggests that greater employment of risk assessments can play a central role in brand management after registration. For example, one core question in studying risk assessments in the law is whether legal decisions can be effectively mapped onto mathematical relationships. However, the process by which human decision-makers give effect to legal rules is inherently a black box. The 2(d) "likelihood of confusion" test reflects the way that law typically creates somewhat nebulous rules. These rules only become effective because human beings (judges, bureaucrats, etc.) interpret them and create standards for how they should be applied. Giving explicit written reasons for decisions is one way that decision-makers can communicate how they arrived at decision.²⁸⁶

Importantly, we echo Rai's central point that transparency and explicability are not necessarily the same thing in the intellectual property context.²⁸⁷ Explicability is an elusive goal in these sorts of agency decisions because human decision-making is inherently a black box. Similarly, machine learning models may also suffer from this lack of explicability.²⁸⁸ In the context of a 2(d) denial of a trademark application, it may be impossible to truly explain how either a trademark officer or a machine learning models making determinations about likelihood of confusion.

However, as we show, not all hope is lost because one need not understand precisely why a potential mark will be rejected as a 2(d) violation in order to make decisions. Simple diagnostic tools can provide insights into how decisions are being made. In our case, we evaluate trademark search engines that deploy AI to power their results and find that they in general reduce search costs for potential users. In doing so, we demonstrate that one way forward in studying risk assessments in the law is to evaluate the outputs of AI models. We specifically focus on search results in trademark search engines, but this general framework could be applied broadly across various domains.

285. *Id.* at 10.

286. *See generally* Frederick Schauer, *Giving Reasons*, 47 STAN. L. REV. 633 (1995).

287. *See* Rai, *supra* note 3.

288. *See generally* Selbst et al., *supra* note 281.

Here, the employment of predictive analytics can also help conserve private resources spent on enforcement. To illustrate, at least two search engines use percentage-based scores to assess the risk that these marks, if selected, would cause legal concern.²⁸⁹ If a firm realizes that a mark with a “very high risk” score has been approved by the USPTO, that information will allow it to prioritize taking legal action against the holder of the conflicting mark, rather than wasting resources pursuing marks that are not especially damaging. Alternatively, a firm that selects a mark with a “very high risk” score faces a high level of vulnerability due to the likelihood of a legal challenge to the mark’s selection.

Figures 22 and 23 show what risk assessments look like in the trademark search context. Figure 22 shows an example report from TrademarkNow,²⁹⁰ and the numerical figures on the left indicate a mark’s riskiness of running into a likelihood-of-confusion denial. Figure 23 shows the distribution of risk scores from Markify.²⁹¹ A user can use these services to prioritize their search results, and evaluate their registration strategy in light of risk assessment scores. From a user’s perspective, one can easily focus on “high” risk results and determine whether to proceed on that basis, while paying less attention to “medium” and “low” risk results. This sort of prioritization is important because the heart of AI-driven trademark search is to reduce the human effort needed to assess likelihood of confusion, and instead focus on other parts of the trademark application process. Because there is a huge supply of potentially conflicting trademarks, the effort required to make a determination about each potential conflict can add up quickly. As we showed earlier, any given searched mark could be expected to return at least ten potential conflicts, and sometimes in excess of two hundred. An AI-generated risk score removes much of this guess work, and would be especially helpful for edge cases. The user can focus on the “high” risk results and tailor their application to avoid conflict with these results. Without wasting time and effort on marks that would be unlikely

289. Both TrademarkNow and Markify provides these assessments. See, e.g., *Unlimited Trademark Screening & Analysis with ProSearch™*, MARKIFY, <https://www.markify.com/services/prosearch-temp.html> (last visited Jan. 23, 2021) (discussing its metrics for “statistical risk analysis”). TrademarkNow’s product description says that its services allow a user to “a clear picture of risk across all regions of interest in seconds and review your clearance search results ranked and analyzed in order of threat.” *Clearance Search – NameCheck™*, TRADEMARKNOW, <https://www.trademarknow.com/products/namecheck> (last visited Jan. 23, 2021).

290. This image is taken from TrademarkNow’s demo page: <https://www.trademarknow.com/name-check-video>.

291. These are drawn from Markify’s Comprehensive Reports rather than the knockout searches we used earlier.

to cause problems anyway, the user would save a potentially enormous amount of time and costs associated with hiring a trademark attorney.

Figure 22: Sample Risk Scores from TrademarkNow's Platform

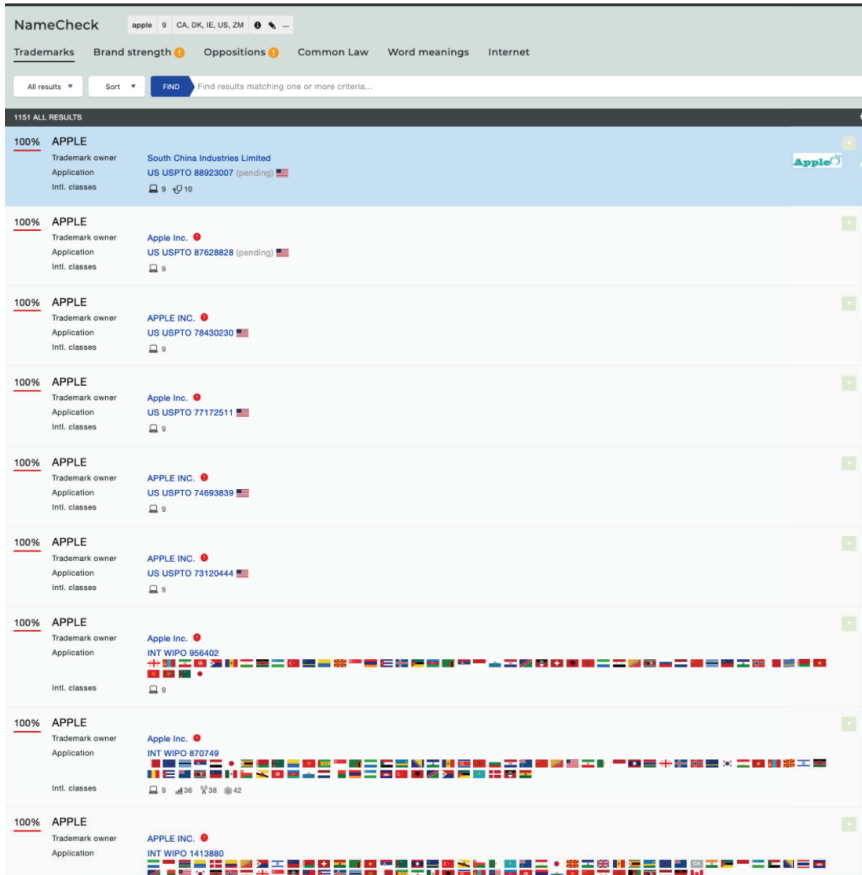



Figure 23: Risk Level Distributions in Markify Dataset



Another way we see the various ways by which underlying AI may work is by looking at how each search engine deals with similarity. Again, we do not know the exact mechanics of how each search engine defines similarity or the thresholds that each chooses when optimizing information retrieval. However, we do have the outputs and can diagnose how well those outputs fit our predefined metrics. In particular, we can use Levenshtein distance²⁹² to analyze the results produced by each search engine. A Levenshtein distance is calculated between two text strings by looking at the number of edits—additions, subtractions, substitutions, and deletions—that it takes to get from one string to another. Figure 24 shows the distribution of Levenshtein distances across some of our search engines. A quick look at each search engine’s distributions shows how their underlying algorithms may prioritize different kinds of results. For instance, Corsearch returns relatively few extremely close matches, likely because its algorithm is more focused on phonetic matching. Trademarkia returns a relatively large number of exact or close matches, indicating that it is more concerned with finding obvious candidates.

292. See generally Frederic P. Miller & Agnes F. Vandome, DAMERAU-LEVENSHTEIN DISTANCE: INFORMATION THEORY, COMPUTER SCIENCE, VLADMIR LEVENSHTEIN, STRING METRIC, STRING (COMPUTER SCIENCE), TRANSPOSITION (MATHEMATICS) (John McBrewster, Ed., 2010).

Figure 24: Distribution of Levenshtein Distances





Figure 25: Lollipop Chart of Median Levenshtein Distances



Delving deeper, we can see this relationship even more clearly. Figure 24 shows the median Levenshtein Distance for results, separated by search engine. TESS, Markify, Trademarkia, and TrademarkNow all tend to return results that are fairly close to the searched mark. Corsearch is a clear outlier here, again, because its algorithm is likely prioritizing different kinds of results.

Using these simply defined metrics and plots, we can see how these relatively straightforward tools can be used to understand and diagnose AI systems. In particular, by focusing our attention on the outputs of these search engines, we can perform apples-to-apples comparisons among them to make inferences about how their underlying algorithms work. These inferences can

then enrich our general theory of search costs in the trademark spaces because they suggest that firms look for a variety of attributes within their initial search for a trademark. These attributes are likely directly related to the outputs that we uncovered in this study, giving us insight into how users make decisions about which AI tools to use in their searches and brand management efforts.

Here, it is important to note that given the sheer power of predictive analytics coupled with massive amounts of data storage and retrieval, there is at least some potential for AI to surpass human judgment and performance when it comes to analyzing and integrating a much wider array of variables in its assessments.²⁹³ But this may not always be a good thing, particularly where subjective judgment (or survey evidence) is relied upon in court. In some cases, risk assessments can result in a mechanistic, formalistic prediction of liability. Where AI lacks the human ability to consider context, it may result in a higher, expanded prediction of likelihood of confusion.²⁹⁴ This outcome suggests at first that a greater reliance on AI at the front end in the registration process may actually reduce the incidence of infringement and confusion at the back end (after the mark has entered the market).²⁹⁵ But this may leave out the consumer in the process of determining actual confusion on the back end. In fact, Dev Gangjee has observed, “[t]he reactions of a real-world consumer, so often alluded to in trademark doctrine, may be muted even further as a result.”²⁹⁶

There are other concerns raised by an overreliance on AI in risk assessment strategies. Given the large number of marks that are not in use, but which remain registered or may be unregistered, there is also a risk that assessments may not reflect the reality of the existing marketplace. Here, AI-driven tools may not be able to distinguish between marks that are actually in use from those that are just claimed for use (but not actually in use yet), thereby creating a greater risk of false positives for likelihood of confusion.²⁹⁷ The converse of this is also created by the limited ability of AI to accurately assess other risks beyond infringement. For example, the risk of dilution through blurring or tarnishment or inclusion of common law trademarks in assessments present other risks that produce more false negatives and enable potential free-riding activity.²⁹⁸

293. See Gangjee, *supra* note 6, at 11.

294. See *id.* at 12–13.

295. See *id.* at 13.

296. See *id.*

297. See *id.* at 14.

298. See *id.*

Future studies, of course, could conduct a similar analysis to study these aspects of trademark search engines. One could generate a list of valuable trademarks and run tests on each search engine to determine how well they flag potential conflicts. Again, the framing here is important. Whereas we looked at the economics from the perspective of a registrant, there is also a fascinating world of study to explore from the perspective of a trademark holder, after a trademark has been granted. One core area worth studying further is how AI fits into an emerging divide in trademark law between those who benefit from utilizing an enforcement strategy focused on litigation and those who do not.²⁹⁹ There may be other ways to generate data surrounding new trademark applications or enforcement strategies, and new experiments could lead to novel new insights.

Last, while our empirical study is limited to basic word search marks, there is room to explore all of the ways that AI is transforming the trademark search space in terms of visual marks and logos, as well. As computer vision tools develop, a follow up study could see how well each search engine returns close visual matches. This sort of study would be fascinating because it would present an interesting exploration of how brands protect elements of their logos and how the USPTO thinks about visual similarity.

D. FUTURE WORK

Our study opens up several possibilities for future work on trademark search and artificial intelligence. In particular, we have established a reproducible method for searching trademark applications, and evaluating how well various search engines do on various metrics. Other researchers can expand the set of searches, change the metrics, or analyze new data in different ways.

In particular, one possible extension of our work is using pending trademark applications instead of previous applications that got 2(d) citations. One could scrape new trademark applications, search these names in the search engines, and wait to see which ones are rejected by the USPTO. This type of study effectively achieves what we did with previous 2(d) citations, but with the benefit of evaluating marks that have yet to be reviewed by the USPTO.

Otherwise, a further area of study could be examining whether there are differences between different types of registrants. Although we did not use this information in our analysis, trademark applications also have information about the registrant. Examining whether there are differences in applications

299. On this point, see generally Glynn Lunney, *Two-Tiered Trademarks*, 56 HOUS. L. REV. 295 (2018).

and 2(d) denial rates among different types of registrants could be interesting. For instance, examining the difference between repeat registrants and first-time registrants, companies in different industries, and various other factors could further enrich our understanding of how trademark search engines work.

Finally, we raise questions about the interplay between AI-powered trademark searches and USPTO trademark-granting activity. This area has been explored theoretically in patent literature already, and we expand this discussion to trademarks. While we provide some preliminary evidence about how trademark search engines work, more work should be done to study the interplay directly and how trademarks evolve over time, if at all.

CONCLUSION

In this paper, we outlined a framework for understanding the economics of trademarks from the perspective of trademark holders, and we examined how AI is rapidly changing the search costs involved with trademark registration and acquisition. We then conducted a novel empirical study that explores how AI is used by trademark search engines, comparing the results from various AI-related private vendors. Our research suggests a greater need for trademark scholars to consider a foundational transformation attributable to AI, where the trademark holder essentially becomes a consumer of trademarks. Such a transformation necessitates a greater attention to the supply of, rather than the demand for, trademarks. Finally, we discussed the implications our findings have for IP law, and the role of AI and search in legal contexts. Going forward, we hope this paper opens up an exploration of the impact that AI will have on trademarks, search costs, and legal administration more broadly.

GENERATING REMBRANDT: ARTIFICIAL
INTELLIGENCE, COPYRIGHT, AND
ACCOUNTABILITY IN THE 3A ERA—THE
HUMAN-LIKE AUTHORS ARE ALREADY
HERE—A NEW MODEL

*Shlomit Yanisky-Ravid**

2017 VISIONARY ARTICLE IN INTELLECTUAL PROPERTY
LAW

2017 MICH. ST. L. REV. 659

ABSTRACT

Artificial intelligence (AI) systems are creative, unpredictable, independent, autonomous, rational, evolving, capable of data collection, communicative, efficient, accurate, and have free choice among alternatives. Similar to humans, AI systems can autonomously create and generate creative works. The use of AI systems in the production of works, either for personal or

* Professor of Law; Yale Law School, Information Society Project (ISP), Fellow; Ono Academic Law School, Israel (OAC), Senior Faculty; Fordham University School of Law, Visiting Professor; The Shalom Comparative Legal Research Center, SCLRC, OAC, Founder and Director. Thanks to Nestor M. Davidson, Joel Reidenberg, Samuel Estreicher, Gideon Parchomosvsky, and Avi Bell as well as to Yale Law School's Information Society Program (ISP) and especially to Jack Balkin, Rebecca Crootof, Michael Fischer, Bonnie Kaplan, and Christina Spiesel. I am grateful to Michele Woods, Michal Svantner, and Victor Lopez Vazquez at the World Intellectual Property Organization (WIPO) in Geneva, Switzerland, to the Swiss Institute of Comparative Law in Lausanne, Switzerland, and in particular to Lukas Heckendorn Urscheler, Alberto Aharonovitz, Ilaria Pretelli, Karen Druckman, and Sadri Saieb for enabling the conduct of this research. I am thankful to the organizers of and participants in the Beyond Intellectual Property Course at Fordham Law School for the important feedback on the research. Thanks also for the Yale Law, ISP and the SCLRC, OAC, Israel; Urbino University, Italy; WIPO, Geneva; and the British Institute of International, Comparative Law (BIICL), London, workshops on the Challenges of Advanced Technology and Artificial Intelligence in the 3A Era, the Day After, for the fruitful discussions in regard to this research. Last but not least, thanks to Samuel Moorhead, Fordham Law School, Intellectual Property Institute, Fellow, and to my gifted research assistants and outstanding legal editors: Alina Lindblom, Iulia Padeanu, Sophie Hogan, Marie Matousek, and Sean Hallisey.

manufacturing purposes, has become common in the 3A era of automated, autonomous, and advanced technology. Despite this progress, there is a deep and common concern in modern society that AI technology will become uncontrollable. There is therefore a call for social and legal tools for controlling AI systems' functions and outcomes.

This Article addresses the questions of the copyrightability of artworks generated by AI systems: ownership and accountability. The Article debates who should enjoy the benefits of copyright protection and who should be responsible for the infringement of rights and damages caused by AI systems that independently produce creative works. Subsequently, this Article presents the AI Multi-Player paradigm, arguing against the imposition of these rights and responsibilities on the AI systems themselves or on the different stakeholders, mainly the programmers who develop such systems.

Most importantly, this Article proposes the adoption of a new model of accountability for works generated by AI systems: the AI Work Made for Hire (WMFH) model, which views the AI system as a creative employee or independent contractor of the user. Under this proposed model, ownership, control, and responsibility would be imposed on the humans or legal entities that use AI systems and enjoy its benefits. This model accurately reflects the human-like features of AI systems; it is justified by the theories behind copyright protection; and it serves as a practical solution to assuage the fears behind AI systems. In addition, this model unveils the powers behind the operation of AI systems; hence, it efficiently imposes accountability on clearly identifiable persons or legal entities. Since AI systems are copyrightable algorithms, this Article reflects on the accountability for AI systems in other legal regimes, such as tort or criminal law and in various industries using these systems.

TABLE OF CONTENTS

INTRODUCTION	661
I. WHAT ARE ARTIFICIAL INTELLIGENCE SYSTEMS?	
HISTORICAL AND TECHNICAL PERSPECTIVES.....	672
A. The Different Kinds of AI Systems: The Old vs. The New and Advanced.....	672
B. How Do Artificial Intelligence Systems Actually Work?	676
C. What Makes Artificial Intelligence Systems Creative? ...	678

Generating Rembrandt

661

II. ACCOUNTABILITY FOR AUTONOMOUS AI SYSTEMS—	
THE COPYRIGHT PERSPECTIVE.....	682
A. Accountability Matters.....	683
B. AI Systems as Independent Legal Entities: The Personhood and Consciousness Approach vs. The Firm Approach	684
1. <i>The Personhood and Consciousness Approach to AI Systems</i>	685
2. <i>The Corporate Approach</i>	687
C. Behind Every Robot There Is a Person: Looking for the Human(s) Behind the Machine.....	689
1. <i>Who Could the Owner Be?</i>	691
2. <i>Distinguishing Between the Rights over Artificial Intelligence Software; the Rights of Works Produced by Automated AI Systems; and the Rights of Programmers</i>	693
3. <i>Other Possible Accountable Entities</i>	697
III. THEORETICAL JUSTIFICATIONS	699
A. Law & Economics.....	700
B. Personality and Labor Theoretical Justifications	706
IV. THE MODEL OF AI—WORK MADE FOR HIRE (WMFH)	707
A. Rethinking the WMFH Legal Doctrine in the Case of AI Systems.....	708
B. WMFH and Works Generated by AI Systems	711
C. The Legal Implications of the AI WMFH Model	713
D. The Drawbacks of Adopting the WMFH Model in Cases of AI Systems.....	714
E. The Advantages of the Proposed AI WMFH Model.....	716
V. U.S. COPYRIGHT LAW IS UNPREPARED TO DEAL WITH AI SYSTEMS.....	718
A. Humans vs. AI Systems as Creators.....	718
B. Eligibility for Copyright Matters.....	720
1. <i>Originality</i>	721
2. <i>Fixed in a Tangible Medium</i>	722
3. <i>Creativity</i>	722
CONCLUSION.....	725

INTRODUCTION

The artist appraises the work, silently judging each stroke of dark ink on the canvas. Determining that the composition is not shaded quite right, the artist decides to switch to an even blacker hue.

Retrieving the brush from the palette, the artist begins to work again, methodically filling the canvas with terse, precise brushstrokes. This is a familiar scene, one that has been playing out in artists' workshops from the medieval classic painters to modern creative artists. This artist, however, is different. It is a robot. Named e-David by its creators at the University of Konstanz in Germany, this robotic artist uses a complex visual optimization algorithm to create paintings.¹ E-David represents merely one step in the ongoing development of the complex, advanced, automated, autonomous, unpredictable, and evolving artificial intelligence (AI) systems that already create original intellectual property works.²

These AI systems are quite different from simple laser printers, which can only reproduce or copy existing works, in a predictable, structural method. E-David, on the other hand, unlike the traditional systems, can produce new drawings in a non-anticipated and creative way.³ E-David does not copy other works, but instead autonomously takes pictures with its camera and draws original paintings from these photographs. Some of these artworks might be entitled to

1. See Oliver Deussen et al., *Feedback-Guide Stroke Placement for a Painting Machine*, in PROC. EIGHTH ANN. SYMP. ON COMPUTATIONAL AESTHETICS IN GRAPHICS, VISUALIZATION & IMAGING 25, 25, 27 (2012) (describing the e-David painting machine, designed to simulate human painting processes, and the methods used by the developers in the Department of Computer and Information Science at the University of Konstanz, Germany). E-David has an arm, five brushes, a camera, a system of optimization via visual feedback, and a system of optimization strategy. See *id.*; see also Jason Falconer, *e-David the Robot Painter Excels in Numerous Styles*, NEW ATLAS (July 17, 2013), <http://newatlas.com/edavid-robot-artist-painter/28310/> [<http://perma.cc/R6DR-44EM>] (arguing that “[t]he line between art and technology isn’t just being blurred, it’s being erased altogether”).

2. Falconer, *supra* note 1 (describing the artworks of e-David as composed of sketches from existing pictures as well as new ones taken with a camera). Relying on existing works might be considered an infringement of the copyright of the original works either directly or as creating derivative works. However, more sophisticated AI systems can create new artworks without copying or infringing copyrights of others. These systems are the focus of this Article. See, e.g., Harold Cohen, *Driving the Creative Machine*, ORCAS CTR., CROSSROADS LECTURE SERIES, 1, 3, 5, 7 (Sept. 2010), www.aaronshome.com/aaron/aaron/publications/orcastalk2s.pdf [<https://perma.cc/5ATB-ALJP>] (describing the AARON machine, which is another machine that creates abstract artworks); see also Harold Cohen, *Fingerpainting for the 21st Century*, AARONS HOME (Feb. 8, 2016), [aaronshome.com/aaron/aaron/publications/8Feb2016Fingerpainting-for-the-21st-Century-with-Figures.pdf](http://www.aaronshome.com/aaron/aaron/publications/8Feb2016Fingerpainting-for-the-21st-Century-with-Figures.pdf) [<https://perma.cc/A2J4-PVSK>] (explaining the techniques and the process of developing the system).

3. See Falconer, *supra* note 1.

Generating Rembrandt

663

copyright protection had humans created them. By using different techniques and an optimization system, e-David makes autonomous and unpredictable decisions about the image it is creating, the shapes and colors, the best way to combine light and shadow, and more.⁴ Even though e-David functions through software created by its programmers, a camera embedded in its complex system allows it to independently take new pictures and generate new creative input as “its own.”

In this Article, I argue that under the “3A era” of automated, autonomous, and advanced technology, sophisticated AI systems and robots turn into talented authors. Indeed, these AI systems already function in the 3A era, generate products and services, make decisions, act, and independently create artworks.

In 2016, nearly 400 years after the death of Rembrandt Harmenszoon van Rijn, the famous Dutch painter, a new Rembrandt, or rather The Next Rembrandt, was unveiled to the world.⁵ The goal of the project was to digitize the painting method of Rembrandt, the human painter.⁶ Once the program “learned” the style of the painter, it would create a new, creative, independent, and original work of art of the genuine Rembrandt.⁷ To ensure its success, the project brought together experts from a variety of fields—engineering, history, art—

4. *Id.* (describing how the software chooses what paint color and brush strokes are needed and how it can make up for inaccuracies in brush strokes and unpredictable paint mixing that occurs on the canvas).

5. Steve Schlackman, *The Next Rembrandt: Who Holds the Copyright in Computer Generated Art*, ART L.J. (Apr. 22, 2016), <http://artlawjournal.com/the-next-rembrandt-who-holds-the-copyright-in-computer-generated-art> [<https://perma.cc/2C2R-EB5N>] (discussing how the first “goal of the project was to discover if an algorithm could . . . produce a physical work of art that would mimic the look of a genuine Rembrandt painting”).

6. *Id.*

7. *Id.*; see also *The Next Rembrandt*, YOUTUBE (Apr. 5, 2016), <https://www.youtube.com/watch?v=IuygOYZ1Ngo&t=3s> [<https://perma.cc/L4PR-NZNC>].

To accomplish this lofty task, the team began with an in-depth study of the proportions and features of the faces in Rembrandt’s works. To master his style, the project team “designed a software system that could understand Rembrandt based on his use of geometry, composition, and painting materials. A facial recognition algorithm identified and classified the most typical geometric patterns used by Rembrandt to paint human features. It then used the learned principles to replicate the style and generate new facial features for our painting.”

Schlackman, *supra* note 5.

and transferred their knowledge into software capable of creating entirely new works of art.⁸

Once a work such as the new Rembrandt painting is created by an AI system, however, policy makers must re-consider the relevancy of the current laws. Can our legal system cope with questions of ownership and responsibility in the 3A era that have never been seen before?⁹ This discussion has deep roots in the copyright regime because AI systems are, ultimately, software algorithms that are regulated under the existing copyright law regime.¹⁰ I argue that one of the main challenges in the near future, the accountability of AI systems, may be solved through the use of copyright lens.¹¹

AI systems and machine learning have already become part of our everyday life. One can already identify AI systems in unexpected regimes, such as: AI doctors,¹² AI therapists,¹³ independent driverless

8. See *The Next Rembrandt*, *supra* note 7.

9. HANOCH DAGAN, PROPERTY: VALUES AND INSTITUTIONS 3-35 (2011) (describing how the ownership of property rights means not only excluding others but also having accountability toward others regarding the right over the property and the use of the property); see also HANOCH DAGAN, RECONSTRUCTING AMERICAN LEGAL REALISM & RETHINKING PRIVATE LAW THEORY 104-28, 161-92 (2013) (disagreeing with the prevailing approach of private law in general and interprets the private law as reflecting horizontal relationships among citizens); Hanoch Dagan, *The Challenges of Private Law: A Research Agenda for an Autonomy-Based Private Law*, in PRIVATE LAW IN THE 21ST CENTURY 67-87 (Kit Barker, Karen Fairweather & Ross Grantham eds., 2017) (advocating for private law as necessary to govern interpersonal relationships).

10. See *Copyright Ownership: Who Owns What?*, STAN. U. LIBR., <https://fairuse.stanford.edu/overview/faqs/copyright-ownership/> [<https://perma.cc/4Y6E-ASJK>] (last visited Jan. 15, 2018).

11. See, e.g., Rebecca Crootof, *War Torts: Accountability for Autonomous Weapons*, 164 U. PA. L. REV. 1347, 1366, 1375-86 (2016) (arguing that because autonomous weapons can independently and unpredictably select and engage targets—causing mass killings and damage—and because there is no individual to blame for reckless behavior, a new legal regime of tort laws should arise in the absence of other existing international tools); see also GABRIEL HALLEVY, WHEN ROBOTS KILL: ARTIFICIAL INTELLIGENCE UNDER CRIMINAL LAW 1-4 (2013) (discussing the accountability of robots for criminal offenses).

12. Jolene Creighton, *AI Saves Woman's Life by Identifying Her Disease When Other Methods (Human) Failed*, FUTURISM (Aug. 5, 2016), <http://futurism.com/ai-saves-womans-life-by-identifying-her-disease-when-other-methods-humans-failed> [<https://perma.cc/8SWR-U9TD>] (“If you needed proof that the age of artificial intelligence is officially upon us, well, look no farther. Reports assert [] that IBM’s artificial intelligence (AI) system, Watson, just saved the life of a Japanese woman by correctly identifying her disease. This is notable because, for

Generating Rembrandt

665

cars,¹⁴ AI lawyers,¹⁵ automated Alternative Dispute Resolution,¹⁶ and automated contracts.¹⁷ AI systems have also significantly influenced many other fields, such as investments,¹⁸ automated weapons,¹⁹

some time, her illness went undetected using conventional methods, and doctors were stumped.”).

13. See Jonathan Amos, *Love Lab Predicts Marital Outcome*, BBC NEWS (Feb. 13, 2004, 9:20 AM), <http://news.bbc.co.uk/2/hi/science/nature/3484981.stm> [<https://perma.cc/ZFU7-Y69K>] (discussing a mathematical model scientists believe can tell which marriages are doomed to end in divorce).

14. See David Szondy, *University of Oxford Develops Low-Cost Self Driving Car System*, NEW ATLAS (Feb. 18, 2013), <http://newatlas.com/oxford-robot-car/26282> [<https://perma.cc/BU7S-6RGY>]; see also Alexandru Budisteanu, *Using Artificial Intelligence to Create a Low Cost Self-driving Car*, BUDISTEANU, <http://budisteanu.net/Download/ISEF%202%20Autonomous%20car%20Doc%20particel.pdf> [<https://perma.cc/Y46J-KSSA>] (last visited Jan. 15, 2018) (discussing how a car that should be able to drive automatically and autonomously in an urban area is achievable). In 2004 road traffic caused 2.5 million deaths worldwide and 50 million injuries—87% of crashes were due solely to driver factors. *Id.* Most of the project’s components of self-driving cars have been done; the system is able to recognize the traffic signs and register them in a common database using Google maps, GPS, and more. *Id.*

15. See, e.g., Jessica Chasmar, *Stanford Student’s Robot Lawyer Has Beaten 160,000 Parking Tickets*, WASH. TIMES (June 29, 2016), <http://www.washingtontimes.com/news/2016/jun/29/joshua-browder-stanford-students-robot-lawyer-has-/> [<https://perma.cc/X9CE-VESH>] (describing a lawyer bot that won 160,000 cases); *This Robot Lawyer Could Help You Get Your Parking Ticket Dismissed*, CBS NEWS (July 21, 2016, 7:05 AM), <http://www.cbsnews.com/news/donotpay-bot-lawyer-helps-dismiss-parking-tickets-joshua-browder> [<https://perma.cc/7REW-JWHC>] (describing Stanford University freshman Joshua Browder and how the robot already saved drivers an estimated \$4-5 million).

16. Chasmar, *supra* note 15.

17. Lauren Henry Scholz, *Algorithmic Contracts*, 20 STAN. TECH L. REV. 128, 133 (2017) (arguing that “[t]he existence of algorithms that must be understood as servants rather than mere tools justifies the creation and analysis of a distinct category called ‘algorithmic contracts,’” and that “[m]achine learning enables sophisticated algorithms to be more similar in function to a human employee with a task to achieve than a tool”).

18. Nizan Geslevich Packin & Yafit Lexv-Aretz, *Big Data and Social Netbanks: Are You Ready To Replace Your Bank?*, 53 HOUS. L. REV. 1211, 1222 (2016) (showing that most financial institutes in North America are using big data analyses and banks are moving toward adopting technologies tools).

19. Rebecca Crotoof, *The Killer Robots Are Here: Legal and Policy Implications*, 36 CARDOZO L. REV. 1837, 1840-43, 1863-71, 1894-1901 (2015) (arguing that AI weapons “systems with varying levels of autonomy . . . have already integrated into the armed forces of numerous states” and calling for defining Autonomous Weapon System and regulating them internationally); see also Roberto Baldwin, *The Robots of War: AI and the Future of Combat*, ENGADGET (Aug. 18, 2016), <https://www.engadget.com/2016/08/18/robots-of-war-ai-and-the-future-of-combat> [<https://perma.cc/NPZ2-66VV>] (arguing that “[t]he future of warfare will be

espionage,²⁰ and even social policymaking.²¹ It is hard to imagine an area of study that has not been influenced by AI systems.

The AI industry has rapidly and consistently become an inevitable part of our present, and it is expected to further develop as the industry is estimated to grow to \$70 billion by 2020.²² Although these systems are set to add substantial value to our world and bring about positive change, there are several drawbacks to these advanced

filled with AI and robots [and] it'll be a world where whoever builds the best artificial intelligent will emerge the victor"); Caitlin Brock, *Where We're Going, We Don't Need Drivers: The Legal Issues and Liability Implications of Automated Vehicle Technology*, 83 UMKC L. REV. 769, 770-73, 787-88 (2015) (arguing that the future of no driver reality is coming and the time to prepare is now); Ray Kurzweil, *The Virtual Thomas Edison*, TIME (Dec. 4, 2000), <http://content.time.com/time/magazine/article/0,9171,90538-2,00.html> [<https://perma.cc/NK3R-29E8>] (discussing issues raised by automated machines and the future of robots).

20. Jasper Hamill, *Eyes in The Sky: CIA Training Artificial Intelligence to Spy on Earth from Space Using Computer Vision*, THE SUN (Aug. 25, 2016, 5:19 PM), <https://www.thesun.co.uk/news/1673802/cia-training-artificial-intelligence-to-spy-on-earth-from-space-using-computer-vision> [<https://perma.cc/4Q67-4AWW>].

21. Rob Kling, *Automated Information Systems as Social Resources in Policy Making*, ACM 666, 666 (1978), <http://dl.acm.org/citation.cfm?id=810109> [<https://perma.cc/C8HL-PTBS>]. Automated information systems have been suggested by a number of theorists to aid public policy makers in acquiring more accurate, timely, and relevant information.

This paper reports a study of the uses and impacts of automated systems for policy analysis in 42 municipal governments. Automated analyses are commonly used in municipal governments . . . and are used to support policy suggestions which are often implemented. Automated systems in these settings serve in both educational and political roles.

See id. But see Jack M. Balkin, *The Three Laws of Robotics in the Age of Big Data*, 78 OHIO ST. L.J. (forthcoming 2017) (manuscript at 18-27) (arguing that a characteristic feature of the Algorithmic Society is that new technologies permit both public and private organizations to govern large populations. Behind robots, artificial intelligence agents, and algorithms are governments and businesses organized and staffed by human beings that exercise power over other human beings mediated through new technologies; therefore it is important to keep three rules: good faith; private owners' fiduciary to the public; and transparency).

22. *See Tech CEOs Declare This the Era of Artificial Intelligence*, FORTUNE (June 3, 2016), <http://fortune.com/2016/06/03/tech-ceos-artificial-intelligence> [<https://perma.cc/K5KK-69C4>] (discussing how "[t]ech companies are diving into AI analytics research, an industry that will grow to \$70 billion by 2020 from just \$8.2 billion in 2013" and that "[a]rtificial intelligence and machine learning will create computers so sophisticated and godlike that humans will need to implant 'neural laces' in their brains to keep up").

Generating Rembrandt

667

systems. Some of these drawbacks include, among other hazards, damage, suffering, and, most significantly, the loss of control. The main legal challenge remains: Who owns the products generated by AI systems and who is responsible for the possibly negative outcomes stemming from them?

Although diverse solutions have been proposed for dealing with the important issue of accountability for the works generated by autonomous AI systems, no one has yet seriously considered the solutions hidden within the paradigms embedded in the law of copyright. This Article proposes a new solution for dealing with the primary struggle regarding accountability of AI systems based on the copyright regime. The Article will address the fundamental intersection of AI systems and intellectual property laws. The Article proposes a solution taken from the copyright domain, one that might further influence the discussion of accountability for other products, such as autonomous cars and weapons, the drug industry, communication, and more. This relationship and the proposed solution (the new Model) have not been extensively discussed in the current literature. In an attempt to fill this gap in the literature, this Article will focus solely on the copyright regime.

Are creative systems such as e-David and The Next Rembrandt a unique phenomenon within the copyright arena? Not at all. Interestingly, the AI industry has not skipped the creative and innovative production of intellectual property and especially copyrightable works. Paintings generated by AI systems are displayed in exhibitions worldwide.²³ A scene in *Ex-Machina*, an independent thriller illustrating the power of AI, raises important questions of copyright law. In the movie, Ava, a humanoid robot, gives Caleb a drawing she has created for him as a gift to gain his love and his trust.²⁴ Ava's creative work was not a reproduction; it was an original piece of art that meets all the criteria for copyright protection, with the exception that it was created by an AI system.²⁵

23. See for example the exhibition of Trevor Paglen, *A Study of Invisible Images* (Sept. 8–Oct. 21, 2017 at Metro Picture, Gallery, NYC, USA), <http://www.metropictures.com/exhibitions/trevor-paglen4/> [<https://perma.cc/3NCW-B96F>] (showing the spectacular exhibition of paintings made by one AI system—the Generator/the painter—with the sophisticated feedback of another AI system—the Discriminator/the trainer—after exchanging millions of examples between these two AI systems. This technique named Generative Adversarial Networks (GANs) uses AI algorithms by implementing two neural networks used in unsupervised machine learning contesting each other).

24. *EX MACHINA* (Universal Pictures International 2015).

25. *Id.*

However, copyright works created by AI systems are no longer just the stuff of science fiction movies.²⁶ Automated machines, or AI-like systems, are already producing original works in almost every copyrightable medium, such as music,²⁷ poetry,²⁸ literature,²⁹ news,³⁰ and many others.³¹ Indeed, today it is almost impossible to imagine any kind of art developed without using at least some digital means. Eventually, automated systems will replace both creators and producers of numerous types of works, products, and services.

Following these latest developments, the legal challenge in the 3A era is to decide who owns the copyright once an automated, autonomous, and advanced machine, or any form of AI system, generates original and creative works independently of the humans who created the AI system itself.³² Subsequently, it is unclear who is entitled to the licensing rights to the product, who is entitled to the

26. Brad Merrill, *It's Happening: Robots May Be the Creative Artists of the Future*, MAKE USE OF (Dec. 17, 2014), <http://www.makeuseof.com/tag/happening-robots-may-creative-artists-future/> [https://perma.cc/8AY7-NPDA].

27. William T. Ralston, *Copyright in Computer-Composed Music: HAL Meets Handel*, 52 J. COPYRIGHT SOC'Y U.S.A. 281, 306 (2005) ("The question of whether machine-generated expression is a proper subject for copyright has been, and probably will continue to be, a subject of continued debate.").

28. Samuel Gibbs, *Google AI Project Writes Poetry Which Could Make Vagon Proud*, THE GUARDIAN (May 17, 2016, 7:01 AM), <https://www.theguardian.com/technology/2016/may/17/googles-ai-write-poetry-stark-dramatic-vogons> [https://perma.cc/9938-ZASR] (discussing how Google, Stanford University, and others are working on an artificial intelligent program that will write poems after exposing the program to novels).

29. Alison Flood, *Computer Programmed to Write Its Own Fables*, THE GUARDIAN (Aug. 6, 2014, 4:09 AM), <http://www.theguardian.com/books/2014/aug/06/computer-programmed-to-write-fables-moral-storytelling-system> [https://perma.cc/6FAC-RL9A] (discussing how a computer can write new and creative stories).

30. For more examples, see Lin Weeks, *Media Law and Copyright Implications of Automated Journalism*, 4 N.Y.U. J. INTELL. PROP. & ENT. L. 67, 87 (2014) (bringing examples of news created by machines and leave the questions regarding copyright issues open); Steve Lohr, *In Case You Wondered, a Real Human Wrote This Column*, N.Y. TIMES (Sept. 10, 2011), <http://www.nytimes.com/2011/09/11/business/computer-generated-articles-are-gaining-traction.html> [https://perma.cc/5SWH-M4RC].

31. Peter Kugel, *Artificial Intelligence and Visual Art*, 14 LEONARDO 137, 137-39 (1981).

32. See Shlomit Yanisky-Ravid & Luis Antonio Velez-Hernandez, *Copyrightability of Artworks Produced by Creative Robots, Driven by Artificial Intelligence Systems and the Originality Requirement: The Formality-Objective Model*, 19 MINN. J.L. SCI. & TECH. 1, 7-8 (arguing that robots that create unique artworks challenge the concept of originality within copyright law and recommending the adoption of a more formal and objective approach).

Generating Rembrandt

669

royalties, and who bears responsibility for copyright infringement and protecting rights from infringements by others (humans or otherwise).³³ Another challenge entails figuring out who is entitled to the moral right,³⁴ if anyone should be at all.³⁵ Should this one role-player take it all or are many different stakeholders targeted?

Take, for example, The Next Rembrandt project. Approximately 350 paintings were analyzed and over 150 gigabytes of digitally rendered graphics were collected to provide the proper instruction set to produce the textures and layers necessary for The Next Rembrandt to have the painterly presence of an original work by the old master.³⁶ Given the hard work involved, the number of people required, and the large monetary investment, one must wonder who bears the responsibility and accountability for these *new* works generated by the AI system? Assuming the owner of the works (which differs from the owner of the AI system) is the most efficient entity to impose accountability on, who should be considered the owner?³⁷ And which legal rights could he or she assert?³⁸

This development re-imagines the whole concept of art and artists, and as such, it has resulted in the need to re-create the legal regime that governs art, especially artworks produced by AI systems.³⁹ Intellectual property in general, and more specifically copyright laws, have become one of the most interesting, challenging, and contrasting fields demonstrating the unique features

33. *See id.* at 6.

34. Xiyin Tang, *The Artist as Brand: Toward a Trademark Conception of Moral Rights*, 122 YALE L.J. 218, 224 (2012); *see generally* Shlomit Yanisky-Ravid, *Rethinking Employees' Intellectual Property Moral Rights: A New Model*, in INTELLECTUAL PROPERTY: INTERDISCIPLINARY PERSPECTIVE (Miriam Bitton & Lior Zemer eds., 2012).

35. *See* Yanisky-Ravid & Velez-Hernandez, *supra* note 32, at 6-7.

36. Schlackman, *supra* note 5. *See* Amanda Levendowski, *How Copyright Law Creates Biased Artificial Intelligence* 3 (Mar. 16, 2017) (unpublished manuscript) (on file with author) (arguing AI is biased because AI needs vast amounts of good data, which is protected by copyright laws that only wealthy entities can afford).

37. *See* sources cited *supra* note 9.

38. *See* Shlomit Yanisky-Ravid & Xiaoqiong (Jackie) Liu, *When Artificial Intelligence Produce Inventions: The 3A Era and an Alternative Model for Patent Law*, CARDOZO L. REV. (forthcoming 2018) (manuscript at 19-22) (arguing that the Multi-Players Model of AI systems places hurdles on entitling one human as the inventor in the case of AI systems produce inventions).

39. For more examples, *see* Yanisky-Ravid & Velez-Hernandez, *supra* note 32, at 13-14.

of advanced technology systems. AI systems can be characterized as creative, unpredictable, independent and autonomous, rational, evolving, capable of data collection and communication, efficient and accurate, and capable of exercising free choice among alternatives.⁴⁰ AI systems are also confronting the traditional concept of looking for the human author behind the creation because the AI systems themselves may “replace” humans.⁴¹

Traditionally, intellectual property laws, and in particular copyright laws, have been based on human creators, who creatively, originally, and independently create works.⁴² But with the advent of AI systems, there is now the possibility that no human is behind the creative process. Instead, AI systems, as automated, autonomous, and advanced machines, create and produce works independently, unexpectedly, and creatively, with self-determination and an independent choice of what to create and how to create it. Even the wrong outcome, such as infringements of the rights of others or counterfeits, may be achieved independently, with no human to blame.⁴³ This raises the pressing issue of whether the human or the AI system should be entitled to ownership rights. This tension between art, creation, and AI systems is no longer a future concern or the topic of a science fiction movie, which is why it merits discussion.

This Article argues that the traditional laws of copyright are inadequate to cope with the new technology involved in creating artworks. I further argue that products and services independently generated by machines challenge the justifications under IP and copyright laws, which rely on humans to create the works. Copyright laws are simply ill-equipped to accommodate this tech revolution and are therefore unlikely to survive in their current form. In order to address the change in the way art is being created, we must either rethink these laws, give them new meaning, or be ready to replace them.

This Article proposes a few alternative scenarios of the new 3A era in which AI systems are capable of generating independent works. After discussing the drawbacks of these scenarios, I propose adopting a new model based on a broader version of the Work Made

40. *Id.* at 7 (describing the features of AI systems).

41. *Id.* at 7-8.

42. CRAIG JOYCE ET AL., COPYRIGHT LAW 3 (10th ed. 2016).

43. See Crootof, *supra* note 11, at 1349, 1376-81 (stating the same argument in regards to autonomous weapons).

Generating Rembrandt

671

for Hire (WMFH) doctrine.⁴⁴ I propose that AI systems should be seen as the creative employee or self-contractor creators working for or with the user—the firm, human, or other legal entity operating the AI system. On the one hand, this proposal reflects and maintains the human features of the AI system, such as independence, creativity, and intelligence. On the other hand, this proposal ensures that the employer or the user maintain the appropriate rights and duties, which include accountability for the outcomes of the AI system. This may be the best solution to the current problem of a lack of accountability for independent AI systems. Seeing the AI system through the copyright lens will provide new opportunities for imposing ownership and accountability on the known legal entities. Implementing a modified WMFH model may structure a feasible solution in the near future and impose responsibilities on the users who have affinities to the AI systems.

Part I of this Article will provide background on AI systems by discussing the different types of systems and their development over recent years. This Part will describe the features that make AI systems intelligent and creative and thus substitutes for human authors. Part II will address the question of who owns, and who takes the responsibility for, works created by AI systems. This Part presents two options. The first option is to see one of the humans or entities involved in the development of an AI system as the one who bears ownership and accountability for the outcomes of that system. The second option is to see the AI systems themselves as the digital, creative, and autonomous authors and hence the owners and the responsible entities for the works they produce. Part III will consider the various theoretical justifications for intellectual property protection. It will examine whether or not these theories lend any support or justification for these options or, alternatively, for a new option. Part IV will discuss the proposed model of AI systems, the WMFH model, and its implications for AI systems. Part V will discuss how U.S. copyright law is unprepared for the recent developments and challenges of AI systems, focusing primarily on the human authorship principle and extending copyright protection to works generated by automated creative AI systems. After determining that existing law is somewhat irrelevant and outdated, I propose that the AI WMFH model can cure not only the inapplicability of current copyright law to new and advanced AI systems, but can also cure the desire to control these systems as well

44. See *infra* Part III.

as to impose accountability on a legal known entity, such as the user of AI systems. By implementing the proposed model—one that sees AI systems as independent contractors or employees of the users and amending the law to accommodate the AI WMFH model—we can control the users of these systems, thus preventing situations in which the public loses control over the unknown outcomes of the AI systems.

I. WHAT ARE ARTIFICIAL INTELLIGENCE SYSTEMS? HISTORICAL AND TECHNICAL PERSPECTIVES

Before discussing the accountability of AI systems from a copyright perspective, one must address more basic questions: How does an AI system work? What does it mean that the system can autonomously create works? I argue that in order to address questions of accountability for AI systems, one must understand what lies beneath the mysterious concept of AI systems. This Part will clarify how automated AI systems function by focusing on one type of AI system that I have named the “pattern recognition” or “similarities identifier” AI system. This understanding is a fundamental step before further discussion takes place concerning the accountability of AI systems from a copyright perspective.

A. The Different Kinds of AI Systems: The Old vs. The New and Advanced

AI algorithms vary significantly.⁴⁵ A diverse array of AI algorithms has been developed to cover a wide variety of data and problems.⁴⁶ This diversity of learning architectures and algorithms

45. M.I. Jordan & T.M. Mitchell, *Machine Learning: Trends, Perspectives, and Prospects*, 349 SCI. MAG. 255, 255 (2015) (representing candidate programs, such as decision trees, mathematical functions, and general programming languages, and searching through these programs, such as optimization algorithms with well-understood convergence guarantees and evolutionary search methods that evaluate successive generations of randomly mutated programs).

46. *See generally* TREVOR HASTIE, ROBERT TIBSHIRANI & JEROME FRIEDMAN, *THE ELEMENTS OF STATISTICAL LEARNING: DATA MINING, INFERENCE, AND PREDICTION* (2d ed. 2009). *See generally* KEVIN P. MURPHY, *MACHINE LEARNING: A PROBABILISTIC PERSPECTIVE* (2012) (offering “a comprehensive and self-contained introduction to the field of machine learning, based on a unified, probabilistic approach” and stressing a principled, model-based approach often using language of graphical models to specify models in a concise and intuitive

Generating Rembrandt

673

reflects the diverse needs of applications capturing different kinds of mathematical structures and offering different levels of amenability to post-hoc visualization and explanation. It provides varying trade-offs between computational complexity, the amount of data, and performance.⁴⁷

Defining AI systems is not an easy task. There are as many definitions as there are types of AI systems.⁴⁸ John McCarthy, who coined the term “Artificial Intelligence,” did not provide an independent definition, while scholars Stuart Russell and Peter Norvig suggested almost ten different definitions.⁴⁹ Definitions generally vary according to the targeted subject, emphasizing different aspects of AI systems.⁵⁰ Based on its features, AI can be defined as a system capable of performing tasks that would normally require human intelligence, such as recognition, decision-making, creation, learning, evolving, and communicating.⁵¹ AI can also be

way); *Overview*, MIT PRESS, <https://mitpress.mit.edu/books/machine-learning-0> [<https://perma.cc/8K3F-BMCY>] (last visited Jan. 15, 2018).

47. Jordan & Mitchell, *supra* note 45, at 257 (arguing that large-scale deep learning systems have had a major effect in recent years in computer vision and speech recognition).

48. Matthew U. Scherer, *Regulating Artificial Intelligent Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARV. J.L. & TECH. 353, 360 (2016) (describing how, unfortunately, there does not yet appear to be any widely accepted definition of AI even among experts, whose definitions vary widely and focus on myriad of ways AI systems are interconnected with human function—the ability to learn, or consciousness and self-awareness—which are difficult to define).

49. STUART J. RUSSELL & PETER NORVIG, *ARTIFICIAL INTELLIGENCE: A MODERN APPROACH* 2-14, 1034 (3d ed. 2010) (describing definitions include thinking and acting humanly, as well as thinking and acting rationally; the definition is based on human features); *see also* Yanisky-Ravid & Liu, *supra* note 38, at 10-11 (listing different definitions of AI systems); *What Is Artificial Intelligence?*, JOHN MCCARTHY’S HOME PAGE (Nov. 12, 2007), <http://www-formal.stanford.edu/jmc/whatisai/node1.html> [<https://perma.cc/4MF3-KJAH>].

50. RUSSELL & NORVIG, *supra* note 49, at 5-12 (discussing different approaches to AI, such as philosophy, psychology, cognitive math).

51. *Id.* at 14; *see also* MARCUS HUTTER, *UNIVERSAL ARTIFICIAL INTELLIGENCE: SEQUENTIAL DECISIONS BASED ON ALGORITHMIC PROBABILITY* 125-26, 231 (W. Brauer, G. Rozenberg & A. Salomaa eds., 2005) (arguing that AI system is a form of intelligence, as a result of features like creativity, problem solving, pattern recognition, classification, learning, induction, deduction, building analogies, optimization, surviving in an environment, language processing, and knowledge). *Artificial Intelligence*, OXFORD DICTIONARY, https://en.oxforddictionaries.com/definition/artificial_intelligence (last visited Jan. 15, 2018) (“The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.”).

described as an instrument that can make existing solutions more efficient by using all data that is reachable by the AI system. Various contexts, such as medical treatments or chess strategies, also lead to different definitions of AI systems.⁵²

Until recently, the “artificial intelligence” field was dominated by quasi-AI systems called “expert systems,” which mainly used a rules-based decision-making process.⁵³ Put more simply, these systems were not fully autonomous and, therefore, not truly “intelligent.” They lacked the ability to learn and produce unpredictable results because they mostly acted in ways predetermined by their human-created programming.⁵⁴ These systems could not evolve through learning. Consequently, they could not be truly creative because they could only “know” information that a human had placed in their “knowledge base.”⁵⁵ Policy makers, nevertheless, still see these systems as the model of advanced technology. In many machines that create artworks, even though the software has some discretion in how to create the final composition, the scope of that discretion is limited to the operation of programming created by the human inventors.⁵⁶ The significance of

52. See Yanisky-Ravid & Liu, *supra* note 38, at 9 (describing why AI systems are intelligent).

53. Dana S. Rao, *Neural Networks: Here, There, and Everywhere—An Examination of Available Intellectual Property Protection for Neural Networks in Europe and the United States*, 30 GEO. WASH. J. INT’L L. & ECON. 509, 509 (1997) (examining “whether U.S. patent law applies to software-implemented neural networks in light of recent decisions by the U.S. Court of Appeals for the Federal Circuit, [as well as] analyz[ing] whether software networks can receive patent protection in the EC, based on Trade Related Intellectual Property Side (TRIPS) agreements and the Berne Convention, EC directives, Member-State statutes, and Member-State case law”).

54. Arthur R. Miller, *Copyright Protection for Computer Programs, Databases, and Computer-Generated Works: Is Anything New Since CONTU?*, 106 HARV. L. REV. 977, 1038-39 (1993) (addressing the claim that it will eventually be impossible to assimilate computer-generated works into the copyright system because they may have no obvious human author, and concluding not only that the case law contains no persuasive objection to extending copyright protection to these works, but also that such an extension would fulfill the constitutional imperative of promoting progress in these areas).

55. *Id.* (concluding that, despite arguments that incorporating new technologies into the current copyright system will lead to overprotection, the current regime is flexible enough to address concerns).

56. See *e-David. A Painting Process*, UNIVERSITÄT KONSTANZ (Apr. 24, 2017) <https://cms.uni-konstanz.de/informatik/edavid/news> [<https://perma.cc/UX4T-XAAB>] (describing the combination of human input and machine learning involved in the creation of the e-David painting robot).

Generating Rembrandt

675

this from the copyright perspective is that human input is still necessary, not only for a work to be produced, but for it to have any sort of creative content. An expert system has become a tool for human creativity.⁵⁷

Even though this type of quasi-AI system still exists, it does not represent the new standard of today, which is the focus of this study. AI technology has advanced rapidly. After working for decades on creating a new type of AI system, computer researchers have recently succeeded in creating a system that can ultimately have serious ramifications for copyright law.⁵⁸ The current AI systems, functioning intelligently and using learning components autonomously, complicate the discussion. These systems are called “neural networks” because they mimic the function of human brains by absorbing and distributing their information processing capacity to groups of receptors that function like neurons; they find and create connections and similarities within the data they process.⁵⁹ Any one of these units, called “perceptrons,” can “know” whether and how much to react given a particular input; taken together, the system of these responses governs the action of the whole machine.⁶⁰ The difference between a neural network and an expert system is that the former model allows the system to “learn” through trial and error.⁶¹ Given a goal, the system can try random outputs until it successfully performs the desired action and then repeat that response the next time it gets the same or a similar input.⁶² Consequently, a neural network could, like a human, “learn” how to paint, write, or compose and generate a work whose creative content is not the result of any human intervention. At first glance, the human inventor or programmer of such a machine seems to have no more claim to a copyright in such a work than an artist’s mother has to her child’s work, or than a camera manufacturer has to the photos taken by photographers, or than a piano manufacturer has to the melody being created by the musicians while using the instrument. After all,

57. Miller, *supra* note 54, at 980. “A congressional committee has held ‘oversight’ hearings on the subject but has taken no action.” *Id.* at 980 n.7 (citing *Computers and Intellectual Property: Hearings Before the Subcomm. on Courts, Intellectual Prop., and the Admin. of Justice of the House Comm. on the Judiciary*, 101st Cong., 1st & 2d Sess. 1 (1989 & 1990)).

58. Rao, *supra* note 53, at 509.

59. *Id.* (discussing how the developments of neural networks, which allows a system to “learn” information while training, has recently rapidly expanded).

60. *Id.* at 511.

61. *Id.*

62. *Id.* at 509.

neither the inventor and programmer nor the mother nor the manufacturer contributed anything to the creative process except the artist him-, her-, or itself.

Following Scherer's evasive definition of an intelligent system—"machines that are capable of performing tasks that, if performed by a human, would be said to require intelligence"⁶³—one may still ask, what makes the system so intelligent? In other words, how does the system really work?

B. How Do Artificial Intelligence Systems Actually Work?

The process of recognition involves the classification or identification of objects, persons, events, or situations. Research about the human brain promoted the development of one group of algorithms, AI (sometimes named by its learning capability—Machine Learning (ML)), capable of identifying objects or automatically classifying them in a similar way to what we believe and know about human perception and pattern recognition.⁶⁴ One way the AI system functions, among many others, is by following the process of human perception in a few stages.⁶⁵ First, the algorithm is presented with multiple examples and their correct classification (pictures of dogs, faces, signals from the body, or any other data that can be subject to patterns of similarities). Second, the algorithm breaks the data down into "tiny" electronic signals, undetectable by humans, and tries to identify hidden insights, similarities, patterns, and connections—without being explicitly programmed on where to look ("training").⁶⁶ Thus, the patterns and

63. Scherer, *supra* note 48, at 362-64 (arguing for a reform in tort law regulation to cover AI systems liability).

64. See Mauricio Orozco-Alzate & Germán Castellanos-Domínguez, *Nearest Feature Rules and Dissimilarity Representations for Face Recognition Problems*, in FACE RECOGNITION 337, 337-56 (Kresimir Delac & Mislav Grgic eds., 2007); see also Mauricio Orozco-Alzate & César Germán Castellanos-Domínguez, *Comparison of the Nearest Feature Classifiers for Face Recognition*, 17 MACHINE VISION & APPLICATIONS 279, 279 (2006) [hereinafter Orozco-Alzate & Castellanos-Domínguez, *Comparison of the Nearest Feature Classifiers*].

65. See generally Orozco-Alzate & Castellanos-Domínguez, *Comparison of the Nearest Feature Classifiers*, *supra* note 64.

66. Anders Krogh, *What Are Artificial Neural Networks?*, 26 NATURE BIOTECHNOLOGY 195, 195-97 (2008) (describing generally how AI systems work). See generally James J. DiCarlo, Davide Zoccolan & Nicole C. Rust, *How Does The Brain Solve Visual Object Recognition?*, 73 NEURON 415 (2012) (explaining that neuroscientists are providing new clues and constraints about the algorithmic solution). See in practice *Datasets For Machine Learning & Artificial Intelligence*

Generating Rembrandt

677

similarities that the algorithm finds (or creates) may not be clear or completely understood by the programmers, trainers, or those who actively functionalize the system. In fact, “[m]any developers of AI systems now recognize that, for many applications, it can be far easier to train a system by showing it examples of desired input-output behavior than to program it manually by anticipating the desired response for all possible inputs.”⁶⁷ Astonishingly, the trainer can be human or another AI system.⁶⁸ Third, performance improves with experience and evolves with new data to which the system is exposed.⁶⁹ In other words, the system is constantly evolving as a result of new data it has either autonomously found or been inputted with by data providers. For example, if we would like the AI system to create music, we would expose it to many songs or rhythms from different clusters of music, and the AI system would find interconnections unfamiliar even to the programmer. The AI system would keep evolving when exposed to new music in the future and would eventually be able to create new original music independently and without copying other works.⁷⁰ A similar process would take place for writing new stories, painting, creating dances, programming design, programming software, detecting signals in roads, producing new drugs, and even designing AI systems.⁷¹

(AI) Training, CLICKWORKER, <https://www.clickworker.com/machine-learning-ai-artificial-intelligence/> [<https://perma.cc/CAA3-YKEW>] (last visited Jan. 15, 2018).

67. Jordan & Mitchell, *supra* note 45, at 255 (illustrating the widespread nature of the adoption of data-intensive machine-learning methods).

68. *See supra* note 23 and accompanying text.

69. *See, e.g.,* Larry Hardesty, *Artificial-Intelligence System Surfs Web to Improve Its Performance*, MIT NEWS (Nov. 10, 2016), <http://news.mit.edu/2016/artificial-intelligence-system-surfs-web-improve-performance-1110/> [<https://perma.cc/2DJK-JKBT>].

70. William Hochberg, *When Robots Write Songs*, THE ATLANTIC (Aug. 7, 2014), <https://www.theatlantic.com/entertainment/archive/2014/08/computers-that-compose/374916> [<https://perma.cc/SMQ6-LCDY>]. EMI is a software program that, although not intelligent, has produced aesthetically convincing new music. Intelligence seeks survival by the exercise of power over a surrounding environment. In composition, intelligence equals decision making. Every composition results from the selection of a finite set of constraints to operate on selected materials; even the most intuitive decision remains itself a decision, and consequently, a product of constraints. *See* Patrício da Silva, *David Cope and Experiments in Musical Intelligence*, SPECTRUM PRESS 1-36 (2003), <http://eclass.uoa.gr/modules/document/file.php/MUSIC124/%CE%94%CE%B9%CE%B1%CE%BB%CE%AD%CE%BE%CE%B5%CE%B9%CF%82/da-silva-david-cope-and-emi.pdf> [<https://perma.cc/Q8KG-FQRB>].

71. *See also* Rana el Kaliouby, *This App Knows How You Feel – From the Look on Your Face*, TED (2015), https://www.ted.com/talks/rana_el_kaliouby_this

We have already been caught unprepared by the latest developments. Traditional intellectual property laws have become irrelevant for new AI systems. Other fields, such as tort and criminal law, may also be unable to solve the emerging issues. Furthermore, the developments are proceeding rapidly. We have to cope not just with existing automated AI systems that create independent, creative, and original artworks, but we also have to be ready for the next generation of AI that will be capable of unsupervised learning, a paradigm in machine-learning research that uses random methods in unexpected and dangerous ways.⁷²

C. What Makes Artificial Intelligence Systems Creative?

Over the past two decades, AI has grown from a laboratory curiosity to a practical technology. It has emerged as an important tool in developing practical software for computer vision; speech recognition; natural language processing; and creating artworks, inventions, and other applications.⁷³ To understand the challenges posed by AI-created artworks, it is important to understand how automated AI systems produce new and creative works, which would have been copyrightable had humans created them.⁷⁴

I identify ten features of AI systems' algorithms that are important to the discussion of accountability of AI systems based on the copyright discourse.⁷⁵ AI systems can be embedded with all or some of these features, all of which are interrelated and partially overlapping. By using these ten features, AI systems are designed to independently create works of useful art.⁷⁶

app_knows_how_you_feel_from_the_look_on_your_face [https://perma.cc/FY39-29AN].

72. HASTIE, TIBSHIRANI & FRIEDMAN, *supra* note 46, at 18-22 (stressing a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way).

73. Yanisky-Ravid & Liu, *supra* note 38, at 2.

74. *Id.*

75. See HALLEVY, *supra* note 11, at 175 (discussing five different attributes that one would expect an intelligent entity to have—communication, internal knowledge, external knowledge, goal-driven behavior, creativity); see also Yanisky-Ravid & Velez-Hernandez, *supra* note 32, at 2 (proposing the adoption of the objective approach to copyright, which enables copyrightability of works produced by creative robots).

76. Jason D. Lohr, *Managing Patent Rights in the Age of Artificial Intelligence*, LEGALTECH NEWS (Aug. 18, 2016), <http://www.legaltechnews.com/id=1202765385194/Managing-Patent-Rights-in-the-?slreturn=20160819081749> [https://perma.cc/6BTC-9DLR] (arguing that much of the AI in the use today is

Generating Rembrandt

679

(1) Creativity. AI systems are capable of more than just copying other works from accessible sources. They operate as creative devices capable of creating entirely new and original works.⁷⁷ This feature is crucial in the intellectual property realm and in particular when discussing copyrightable artworks.

(2) Autonomous and independent.⁷⁸ A device is independent or autonomous if it can accomplish a high-level task on its own, without external intervention.⁷⁹ Such systems may work independently, with minimum human intervention.⁸⁰ In this way, the AI systems are able to replace authors and other creators, to autonomously produce new artworks.⁸¹

(3) Unpredictable and new results. AI systems are based on algorithms capable of incorporating random input, resulting in unpredictable routes to the optimal solution, and hence creating unpredictable works (from the software programmers' point of view).⁸² An AI system can draw a new painting, which, unlike copying an existing work, is new and unpredictable. After being exposed to colors, shapes, and techniques that are in the public

referred to as “soft” AI systems, where the AI uses computational intelligence to analyze relevant data and attempt to solve a specific problem).

77. See HUTTER, *supra* note 51, at 2 (mentioning creativity as one of the main features of AI); see also Scherer, *supra* note 48, at 364-65 (describing how AI systems detected breast cancer prognosis by checking cells of supportive tissues through a chess player creative move); HALLEVY, *supra* note 11, at 176 (arguing that an AI system must be creative by finding alternative ways to solve problems).

78. Crotoof, *supra* note 19, at 1854-63 (describing the difficulty of deciding on a definition for autonomous weapons and suggesting a definition based on the AI (weapon) system being able (1) to come to conclusions (2) derived from gathered information and (3) is capable of independently selecting actions (selecting and engaging targets)).

79. Lucy Suchman & Jutta Weber, *Human-Machine Autonomies*, in AUTONOMOUS WEAPON SYSTEMS: LAW, ETHICS, POLICY 75, 76 (Nehal Bhuta et al. eds., 2016).

80. See Terence Davis, *The AI Revolution: Is The Future Finally Now?*, ARN (Apr. 14, 2017), <https://www.arnnet.com.au/article/617707/ai-revolution-future-finally-now> [<https://perma.cc/UX4T-XAAB>] (“What is called AI even today is in fact, the leveraging of machines with minimal – though not zero – human intelligence to solve specific, narrow problems.”).

81. See generally Yanisky-Ravid & Velez-Hernandez, *supra* note 32.

82. See Jonathon Keats, *John Koza Built an Invention Machine*, POPULAR SCI. (Apr 19, 2006), <https://www.popsci.com/scitech/article/2006-04/john-koza-has-built-invention-machine> [<https://perma.cc/3ZB3-79LJ>].

domain, the system can “break” the data into digital components, recombine them, and create new and unexpected artworks.⁸³

(4) Capable of data collection and communication with outside data. A significant feature of an AI system is that it can actively “search” for outside data. For example, e-David might autonomously take pictures of the outside world and draw them into new, original, and creative works. Communication is thus a necessary feature of an AI system.⁸⁴

(5) Learning capability. Based on the data it has gathered, an AI system can continue to process data by receiving feedback and improving the results.⁸⁵

(6) Evolving. As a result of the new input and the AI system’s capacity for continuous processing, the system might constantly find new patterns and similarities and hence change the outcomes. In this sense, the system is constantly evolving. This feature is at the core of AI and data science.⁸⁶

(7) Rational-intelligent system. An intelligent system is one with a rational mechanism capable of perceiving data and deciding which activities or omissions would maximize its probabilities of success in achieving a certain goal.⁸⁷

83. See Lawrence Hunter, *Molecular Biology for Computer Scientists, in* ARTIFICIAL INTELLIGENCE AND MOLECULAR BIOLOGY 1, 12-15 (Lawrence Hunter ed., 1993) (describing how similarities enable the composition of cells by its parts as membranes, proteins and other parts by AI systems).

84. See generally *id.*; see also Deussen et al., *supra* note 1, at 1 (discussing how, as part of the Rembrandt project, the robot had a camera that kept on photographing); Matthew Field, *Facebook Shuts Down a Robots After They Invented Their Own Language*, THE TELEGRAPH (Aug. 1, 2017, 10:21 AM), <http://www.telegraph.co.uk/technology/2017/08/01/facebook-shuts-robots-invent-language/> [<https://perma.cc/8FRE-67VZ>] (the chatbots were meant to learn how to negotiate by mimicking human trading and bartering; however, when the social network paired two of the programs, nicknamed Alice and Bob, to trade against each other, they started to develop their own bizarre form of communication that the researchers could not understand).

85. RUSSELL & NORVIG, *supra* note 49, at 928-69 (explaining the process of perception of AI systems, in which the systems are being connected to the raw world, including image formation, color, edge detection, texture, segmentation of images, objects recognition, reconstructing the 3D world, and motions).

86. Jordan & Mitchell, *supra* note 45, at 255, 257 (discussing recent progress in machine learning and illustrating the wide-spread nature of the adoption of data-intensive machine-learning methods).

87. RUSSELL & NORVIG, *supra* note 49, at 3-4, 27, 34-56, 973-85 (describing AI systems as being capable of taking “rational” actions based on environmental input); see also HUTTER, *supra* note 51, at 2, 125-26, 231 (discussing how AI systems can solve problems by using features such as learning, induction,

Generating Rembrandt

681

(8) Efficiency. AI systems are capable of accurately, efficiently, and rapidly processing vast volumes of data—well beyond the ability of the human brain.⁸⁸

(9) “Free choice.” AI systems are able to choose between alternatives in order to arrive at the best outcome.⁸⁹ E-David, for example, chooses between lights, colors, and shapes while drawing.⁹⁰

(10) Goal oriented. AI systems function according to goals such as creating, drawing, writing stories or news, or composing melodies or poems.⁹¹

AI systems that create artworks incorporate, to a certain level, all of these ten features. Once we understand these features, and that the AI systems create outcomes independently and autonomously, we realize that the rights available under copyright laws cannot be afforded only to human authors, and thus, the traditional copyright laws may be inapplicable.⁹² As technology advances, AI systems have become increasingly capable of mimicking part of the functions that we once considered intrinsic to the human mind’s creativity. AI

deduction, building analogies and optimization, as well as using knowledge); DAVID L. POOLE & ALAN K. MACKWORTH, *ARTIFICIAL INTELLIGENCE: FOUNDATION OF COMPUTATIONAL AGENTS* 71, 283-334, 597-611 (2010) (describing AI systems as agents of cognitive skills such as: problem solving, searching for data, learning and evolving, rational planning, and more).

88. GEORGE F. LUGER, *ARTIFICIAL INTELLIGENCE: STRUCTURES AND STRATEGIES FOR COMPLEX PROBLEM SOLVING* 26 (6th ed. 2016) (arguing that AI can refer to all programming techniques trying to solve problems more efficiently than algorithmic solutions and can operate close to the intelligence of human behavior); Woodrow Hartzog et al., *Inefficiently Automated Law Enforcement*, 2015 MICH. ST. L. REV. 1763, 1765-67, 1793-95 (arguing that automated machines are more efficient than human but this is a risky factor and that law enforcement of automated machine should preserve inefficiency for ethical reasons).

89. Scherer, *supra* note 48, at 361-62 (arguing that even when AI systems act rationally, they can still pose public risk—killing efficiently, for example).

90. *See supra* note 1 and accompanying text.

91. *See supra* note 66.

92. RUSSELL & NORVIG, *supra* note 49, at 4-7. The discourse about AI systems includes controversial arguments about the philosophy regarding AI systems. For example, can machines perceive and understand (i.e., can they pass the Chinese test)? Are human intelligence and machine intelligence the same (i.e., can they pass the Turing test)? What is intelligence? What does it mean that a machine think or act rationally; can a machine be self-aware? Can a machine be original or creative? *Id.* However, one must also be aware of the “Eliza Effect.” *See* ROBERT TRAPPL, PAOLO PETTA & SABINE PAYR, *EMOTIONS IN HUMAN ARTIFACTS* 353 (Robert Trapp, Paolo Petta & Sabine Payr eds., 2002) (describing the “Eliza Effect” as the tendency for people to treat machines or programs that are responsive as having more intelligence than they really do, as having human traits, and finding analogies between human behaviors and computer behaviors).

systems will be able to improve specific human skills, not only in terms of accuracy or capacity to process vast amounts of data, but also in terms of creativity, autonomy, novelty, and other features necessary for establishing copyrightable works. Moreover, autonomous AI systems will be able to develop new artworks without significant guidance or instructions from humans.⁹³

Generally, the human or entity behind the process is at the forefront of legal discussions. This Article calls for a different solution, one from an alternate point of view—the intellectual property and copyright laws at stake in this area. The inquiry begins with considering whether AI systems may own the products they produce. While this Article agrees that understanding the human-like features of AI may lead to the conclusion that an artwork being generated by an AI system might belong to the AI system, unlike other scholars, this Article argues that the traditional copyright laws may be irrelevant and inapplicable to these situations and that either modifications or other legal tools should replace them.⁹⁴ The next Part will begin by addressing the discourse of ownership and accountability for AI systems producing original works.

II. ACCOUNTABILITY FOR AUTONOMOUS AI SYSTEMS—THE COPYRIGHT PERSPECTIVE

AI systems are commonly used to generate works for personal or industrial goals. Who should benefit from the works being produced by the AI systems? Who should bear responsibility when something goes wrong? In other words, who is entitled to the rights? Who should be accountable when AI systems infringe on third parties' rights or counterfeit existing works? Should it be the programmers, the trainers, the users, or, perhaps, the AI systems themselves?

93. See Ryan Abbott, *I Think, Therefore I Invent: Creative Computers and the Future of Patent Law*, 57 B.C. L. REV. 1079, 1080 (2016) (stating that AI systems and computers are already generating patentable inventions and arguing that AI should receive patent rights in its inventions); see also Lohr, *supra* note 76 (discussing how AI systems will be able to operate without significant guidance or instruction and to develop new products and processes).

94. See Abbott, *supra* note 93, at 1080-81 (stating that AI systems own the IP rights).

Generating Rembrandt

683

A. Accountability Matters

Advanced technologies, such as AI systems, are forcing us as a society to face new ethical and legal challenges and to rethink basic concepts such as ownership and accountability. Scholars have not yet deeply discussed the notion of copyright accountability for infringements involving AI systems, even though AI systems are themselves copyrightable.

According to scholars such as Hanoch Dagan, Michael Heller, and others, ownership of property rights (applicable also to intellectual property rights) is not merely a question of benefits arising from the right to exclude others from enjoying, using, or licensing the objects.⁹⁵ It is also a question of accountability for using it with consideration for other humans' and entities' rights. Moreover, ownership may also entail rights of others to enjoy the property.⁹⁶ This is also true when discussing AI systems. Adopting this accountability for property rights approach of Dagan and Heller into the discussion on intellectual property rights, in regard to works generated by AI systems, allows us to bind together the benefits and accountabilities of ownership.

The main risk we face today and in the near future is that of losing control over the operation of AI systems.⁹⁷ Moreover, we risk losing control not only of one AI system, but also two or more AI systems acting in concert "behind our backs." Therefore, I have decided to focus on accountability for works generated by AI systems⁹⁸ as AI systems threaten all social and legal regimes.⁹⁹

95. Hanoch Dagan & Michael A. Heller, *The Liberal Commons*, 110 YALE L.J. 549, 559-60 (2001) (seeing ownership of property as accountability for others); Hanoch Dagan, *Pluralism and Perfectionism in Private Law*, 112 COLUM. L. REV. 1409, 1421-22, 1438-39 (2012) (property law regimes cannot be based on the right of exclusion alone; rather, they must be based on human relationships).

96. See Schlackman, *supra* note 5.

97. See *Research Priorities for Robust and Beneficial Artificial Intelligence*, FUTURE OF LIFE INST., <https://futureoflife.org/ai-open-letter> [<https://perma.cc/8FL8-UP6Q>] (last visited Jan. 15, 2018).

98. Ariel Ezrachi & Maurice E. Stucke, *Artificial Intelligence & Collusion: When Computers Inhibit Competition* 37-38 (Oxford Legal Studies Research Paper No. 18; Univ. of Tenn. Legal Studies Research Paper No. 267, 2015). See also Field, *supra* note 84.

99. ARIEL EZRACHI & MAURICE E. STUCKE, *VIRTUAL COMPETITION: THE PROMISE AND PERILS OF THE ALGORITHM-DRIVEN ECONOMY* 56-82, 85-144, 147-202 (2016) (describing how consumers reap many benefits from online shopping and how the sophisticated algorithms behind online retail are changing the nature of market competition, including in negative ways. The authors describe one danger as

Professor Jack Balkin describes several problems of AI systems.¹⁰⁰ The first problem entails the distribution of rights and responsibilities among human beings when non-human agents create benefits, like artistic works, or cause harms, like physical injuries.¹⁰¹ The difficulty arises from the fact that the behavior of robotic and AI systems is “emergent,” meaning their actions may be unpredictable or unconstrained by human expectations.¹⁰² Robotics and AI thus feature emergent behaviors that escape human planning and expectations.¹⁰³ Balkin further cautions that we should not consider all features of a technology to be essential without first considering how the technology is used in society.¹⁰⁴ It would thus be unhelpful to codify certain features as “essential” because they may in reality be transient features arising from current uses and social trends.¹⁰⁵

B. AI Systems as Independent Legal Entities: The Personhood and Consciousness Approach vs. The Firm Approach

Many scholars have recently adopted the idea that autonomy, creativity, and spontaneous evolution of AI systems leads to the recognition of AI systems (and robot embedded systems) as independent legal entities entitled to legal and commercial rights and duties.¹⁰⁶ In other words, scholars argue that the AI system is an

being computers colluding with one another. They describe a second danger as behavioral discrimination based on companies tracking and profiling consumers to get them to buy goods at the highest price they are willing to pay. The authors posit a third danger as the “frenemy” relationship between super-platforms and independent app developers. They caution that data-driven monopolies dictate the flow of personal data and determine who gets to exploit potential buyers); Crootof, *supra* note 19, at 1842-43 (describing the threat of tort war over autonomous weapons).

100. See Jack M. Balkin, *The Path of Robotics Law*, 6 CALIF. L. REV. CIR. 45, 45 (2015).

101. See *id.* at 46, 48-49.

102. See *id.* at 45-46.

103. *Id.* at 46 (arguing that robotics and AI raise the “substitution effect,” meaning people will substitute robots and AI agents for living things but only in certain ways and only for certain purposes. Balkin argues this substitution is likely to be incomplete, contextual, unstable, and often opportunistic).

104. See *id.* at 45.

105. *Id.* (contending that innovation in technology is not just about tools and techniques, but also economic, social and legal relations, which in turn affects how technologies may change).

106. SAMIR CHOPRA & LAURENCE F. WHITE, A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS 1-3 (2011) (arguing for the legal personhood of an artificial agent that will soon be independent, and discussing the artificial agent as

Generating Rembrandt

685

autonomous legal entity that may, and should, be responsible for the outcome of its own actions or omissions.¹⁰⁷ This conclusion may be based on two alternative premises. First, the defining features of AI systems—intelligence, rationality, independence, and the like—are similar to those of humans; therefore, they should be treated as independent entities with legal rights and duties. Alternatively, AI systems are analogous to firms, which are separate, non-human legal entities capable of possessing legal rights, benefits, and responsibilities.

1. *The Personhood and Consciousness Approach to AI Systems*

Can robots be human persons and hence entitled to legal rights (and duties)? Can Ava, one of the robots in the movie *Ex Machina*, be considered the owner of the copyright in her painting and have the duty to avoid infringing other humans' or robots' rights?¹⁰⁸ Or can only humans be persons?

“Artificial intelligence already exhibits many human characteristics. Given our history of denying rights to certain humans, we should recognize that robots are [like] people and have human rights.”¹⁰⁹ This statement by Harvard Law Professor Glenn

capable of having “knowledge” and decision-making ability); Abbott, *supra* note 93, at 1080 (arguing that artificial intelligence systems should be considered inventors for the purposes of patent law). *See also* JOHN FRANK WEAVER, *ROBOTS ARE PEOPLE TOO: HOW SIRI, GOOGLE CAR, AND ARTIFICIAL INTELLIGENCE WILL FORCE US TO CHANGE OUR LAWS* 1, 3-4 (2014) (arguing that robots are independent entities).

107. Abbott, *supra* note 93, at 1080 (arguing that computers are already generating patentable subject matters qualifying as inventors and overtaking human inventors as primary source of new discoveries and inventions, and therefore, AI should receive patent rights in their inventions). *See also* Colin R. Davies, *An Evolutionary Step in Intellectual Property Rights – Artificial Intelligence and Intellectual Property*, 27 *COMPUTER L. & SECURITY REV.* 601, 617-19 (2011), <http://www.sciencedirect.com/science/article/pii/S0267364911001518> [<https://perma.cc/RR6K-W56M>] (claiming that the systems can be the authors and the inventors).

108. *EX MACHINA*, *supra* note 24.

109. Glenn Cohen, *Should We Grant AI Moral and Legal Personhood?*, *ARTIFICIAL BRAIN* (Sept. 24, 2016), <http://artificialbrain.xyz/should-we-grant-ai-moral-and-legal-personhood> [<https://perma.cc/ELL3-CQRK>]; *see also* Big Think, *A.I. Ethics: Should We Grant Them Moral and Legal Personhood?*, *YOUTUBE* (Sept. 23, 2016), <https://www.youtube.com/watch?v=gvcbOSAKF2M> [<https://perma.cc/6QAR-2W4N>] (discussing the distinction between people and human beings, and suggesting granting more rights to AI systems so that we do not err and find ourselves on the wrong side of history even though, at the heart of the matter, the idea scares a lot of people).

Cohen reflects not only his claim that AI already does much of what human beings can do, but also the reality that the digital software of AI systems, which mimics human intelligence, is already far superior to our own.¹¹⁰ Ongoing developments in natural language and emotion detection suggest that AI will continue its encroachment on the domain of human abilities.

The personhood approach to AI systems sees the systems as capable of experiencing consciousness. The goal of the artificial consciousness approach is to explore the cognitive abilities in robots.¹¹¹ Igor Aleksander suggested more than a dozen principles for artificial consciousness, including conscious and unconscious states, learning, memorizing, prediction, self-awareness, representation of meaning, language, will, instinct, and emotion.¹¹² The aim of artificial consciousness is to define whether and how these and other aspects of consciousness can be synthesized in an engineered artifact such as a digital computer.

By virtue of modeling itself, AI systems have sensations and are able to make decisions freely. This can be regarded as having consciousness.¹¹³ The ability to produce consciousness—the ability to experience things, which is found in humans as well as in AI systems—means the ability to recognize, allocate, organize, and recall cognitive sources. Consciousness occurs when we have a symbol for things. We do not know what taste or smell means for any individual human, but we can recognize it by connecting it to an existing symbol.¹¹⁴ This may also be true for AI systems. This approach of computationalism sees the human brain, essentially, as a

110. See Cohen, *supra* note 109.

111. James A. Reggia, *The Rise of Machine Consciousness: Studying Consciousness with Computational Models*, 44 NEURAL NETWORKS 112, 112-31 (2013) (describing the artificial consciousness approach also known as AC).

112. See generally Igor Aleksander, *Machine Consciousness*, in BLACKWELL COMPANION TO CONSCIOUSNESS (Max Velmans & Susan Schneider eds., 2007).

113. Drew McDermott, *Artificial Intelligence and Consciousness*, in THE CAMBRIDGE HANDBOOK OF CONSCIOUSNESS 117, 140-150 (Philip David Zelazo, Morris Moscovitch & Evan Thompson eds., 2007) (claiming that tests such as the Turing test and the Chinese box test are not necessarily relevant to the computational theory of consciousness. In Turing's test a person tries to distinguish a computer from a person by carrying on typed conversations with the computer. If the person who judges the system thinks the computer is human about 50% of the time, then the computer passes the test and is considered less distinguishable from a human. The Chinese Box test concerns situations where a machine uses inputs to create reasonable and logical outcomes, but does not "understand" how or why those outcomes are the correct responses).

114. *Id.* at 118.

Generating Rembrandt

687

computer.¹¹⁵ Once we establish the concept of an impersonal level of meaning in brains and computers, we can introduce the idea of a self-model, a device that a robot or a person could use to answer questions about how it interacts with the world.¹¹⁶ This idea was introduced by Minsky almost forty years ago, and has since been explored by others.¹¹⁷ Other scholars claim that “consciousness is a property of complex systems that have a particular ‘cause-effect’ repertoire.”¹¹⁸ They interact with the world in ways similar to the way the brain does. “If you were to build a computer that has the same circuitry as the brain, this computer would also have consciousness associated with it. . . . However, the same is not true for digital simulations.”¹¹⁹

This approach sees the AI system as a person and thus as capable of bearing rights and duties. An alternative approach imposes rights and duties on AI systems from a different angle—that of the firm approach.

2. *The Corporate Approach*

The corporation as a legal entity can serve as a legal basis for imposing rights and duties on AI systems. Corporations are legal entities subject to a legal regime, including corporate, labor, and even criminal law.¹²⁰ Therefore, the question relating to AI entities has become: Does the growing intelligence of AI entities subject them, as any other legal entity, to legal social control?¹²¹

115. *Id.*

116. *See generally id.* at 117-150.

117. *See* MARVIN L. MINSKY, SEMANTIC INFORMATION PROCESSING 1 (Marvin Minsky ed., 1968) (discussing multiple experiments that explored intelligent machines nearly four decades ago); Aaron Sloman & Ron Chrisley, *Virtual Machines and Consciousness*, 10 J. CONSCIOUSNESS STUD. 1, 18 (2003).

118. Antonio Regalado, *What it Will Take for Computers to Be Conscious*, MIT TECH. REV. (Oct. 2, 2014), <https://www.technologyreview.com/s/531146/what-it-will-take-for-computers-to-be-conscious> [<https://perma.cc/JPP5-LBSD>].

119. *Id.*

120. *See* STEVEN BOX, POWER, CRIME AND MYSTIFICATION 16-79 (1983); John C. Coffee, Jr., “*No Soul To Damn: No Body To Kick*”: *An Unscandalized Inquiry into the Problem of Corporate Punishment*, 79 MICH. L. REV. 386, 386-87 (1981); Brent Fisse & John Braithwaite, *The Allocation of Responsibility for Corporate Crime: Individualism, Collectivism and Accountability*, 11 SYDNEY L. REV. 468, 469 (1988).

121. *See generally* Bruce G. Buchanan & Thomas E. Headrick, *Some Speculation About Artificial Intelligence and Legal Reasoning*, 23 STAN. L. REV. 40 (1970); E. Donald Elliott, *Holmes and Evolution: Legal Process as Artificial*

There are several consequences to this approach. In Europe, for example, there is a strong movement arguing that robots should pay taxes.¹²² Scholars have also proposed that AI systems should be held liable for any criminal offenses committed by the systems.¹²³

If assessed through the lens of copyright laws, this approach would result in AI systems' ownership of the intellectual property products and processes they generate.¹²⁴ Under this view, the AI system is the protagonist: when it acts autonomously, it is the true creator or producer of the products. In this case, the owner might be the AI system itself. Section 201(a) of the Copyright Act states that "[c]opyright in a work protected under this title vests initially in the author[.]"¹²⁵ The U.S. Supreme Court has explained that, as a general rule, "the author is the party who actually creates the work."¹²⁶ Scholars have also endorsed this position, arguing that the AI system

Intelligence, 13 J. LEGAL STUD. 113 (1984); Antonio A. Martino, *Artificial Intelligence and Law*, 2 INT'L J.L. & INFO. TECH. 154 (1994); L. Thorne McCarty, *Reflections on Taxman: An Experiment in Artificial Intelligence and Legal Reasoning*, 90 HARV. L. REV. 837 (1977); Edwina L. Rissland, *Artificial Intelligence and Law: Stepping Stones to a Model of Legal Reasoning*, 99 YALE L.J. 1957 (1990).

122. Michaela Georgina Lexer & Luisa Scarcella, *The Effects of Artificial Intelligence on Labor Markets – A Critical Analysis of Solution Models from a Tax Law and Social Security Law Perspective* (working manuscript) (on file with the authors) (arguing that robots should pay taxes and describing the European practical approach supporting this idea); see also Chris Weller, *Bill Gates Says Robots That Take Your Job Should Pay Taxes*, BUS. INSIDER (Feb. 17, 2017, 9:57 AM), <http://www.businessinsider.com/bill-gates-robots-pay-taxes-2017-2> [<https://perma.cc/J3DJ-PKKN>] (describing an interview with Bill Gates where he argued that robot tax could finance jobs taking care of elderly people or working with kids in schools, for which needs are unmet and to which humans are particularly well suited).

123. See generally HALLEVY, *supra* note 11 (developing a general and legally sophisticated theory of the criminal liability for AI and robotics).

124. See, e.g., Mark Fischer, *Are Copyrighted Works Only by and for Humans? The Copyright Planet of the Apes and Robots*, DUANE MORRIS BLOG (Aug. 18, 2014), <https://blogs.duanemorris.com/newmedialaw/2014/08/18/are-copyrighted-works-only-by-and-for-humans-the-copyright-planet-of-the-apes-and-robots> [<https://perma.cc/C9Z5-X5AY>] (arguing that the future of copyright may someday be in the hands of non-humans).

125. See 17 U.S.C. § 201(a) (2012) (ownership of copyright).

126. *Community for Creative Non-Violence v. Reid*, 490 U.S. 730, 737 (1989); see also Russ VerSteeg, *Defining "Author" for Purposes of Copyright*, 45 AM. U. L. REV. 1323, 1326 (1996).

Generating Rembrandt

689

should be accountable for the outcome of its own actions or omissions.¹²⁷

Ownership, however, might be a result of a commercial contract and not of copyright laws.¹²⁸ This view of AI systems ultimately considers the AI system to be the owner of its works. Scholars, however, have criticized this view on the grounds that it is an untenable proposition.¹²⁹ Moreover, the length of protection is designed after the life of the creator.¹³⁰ Moral rights, including the entitlement of the author to credit as well as the author's control over changes and modifications to the work, remain unresolved when AI systems generate works.

C. Behind Every Robot There Is a Person: Looking for the Human(s) Behind the Machine

Arthur R. Miller said, “[B]ehind every robot there is a good person.”¹³¹ This phrase, which represents the traditional approach to AI in the U.S. and Europe, supports the default view of programmers as the creators entitled to ownership of the works created by the AI systems they have programmed.¹³² Under this view, ownership and

127. Abbott, *supra* note 93, at 1080 (arguing that computers are already generating patentable subject matters qualifying as inventors and overtaking human inventors as primary source of new discoveries and inventions and therefore, AI should receive patent rights in their inventions).

128. *Id.* at 1115-17.

129. See Annemarie Bridy, *Coding Creativity: Copyright and the Artificially Intelligent Author*, 5 STAN. TECH. L. REV. 1, 26 (2012); Pamela Samuelson, *Allocating Ownership Rights in Computer-Generated Works*, 47 U. PITT. L. REV. 1185, 1226-28 (1985) (arguing that rights should accrue to the user of the program as the best practical solution); Robert Yu, *The Machine Author: What Level of Copyright Is Appropriate for Fully Independent Computer-Generated Works?*, 165 U. PA. L. REV. 1245, 1263-65 (the author suggests the contribution-rights paradox: from a social policy standpoint, entitling the rights to independent computer-generated works is wrong). *But see* Fischer, *supra* note 124 (arguing that the future of copyright may someday be in the hands of non-humans).

130. See, e.g., Roberta Rosenthal Kwall, *Copyright and the Moral Right: Is an American Marriage Possible?*, 38 VAND. L. REV. 1, 87-88 (1985).

131. Miller, *supra* note 54, at 1045.

132. See Robert C. Denicola, *Ex Machina: Copyright Protection for Computer-Generated Works*, 69 RUTGERS U. L. REV. 251, 265, 271, 275 (2016) (stating that a computer user who initiates the creation of computer-generated expression should be recognized as the author and copyright owner of the resulting work); John Frank Weaver, *How Artificial Intelligence Might Monetize Fan Fiction*, SLATE (Dec. 10, 2013, 11:33 AM), http://www.slate.com/blogs/future_tense/2013/

accountability for works generated by AI systems are given to the creators of the AI systems.¹³³ According to this view, the ownership of works generated by AI systems and, hence, the accountability for these works “belong” to the humans (and the entities working on their behalf) involved in the process of developing the AI systems. The human behind the program—usually the programmer—has become an important figure in other fields of law that involve harm and damages resulting from AI systems, such as criminal law or tort law.¹³⁴

This traditional approach is reflected in various European Union laws. For example, the British Copyright, Designs, and Patents Act of 1988 takes the approach that copyright protection is proper for persons responsible for a computer’s creation.¹³⁵ The Act states: “In this Part ‘author’, in relation to a work, means the person who creates it.”¹³⁶ Even the broader approach regarding computers generating artworks is looking for the person behind the creation process. Article 9(3) of the Act says: “(3) In the case of a literary, dramatic, musical or artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken.”¹³⁷

The U.S. also holds this attitude, as reflected by the National Commission on New Technological Uses of Copyrighted Works (CONTU), which was created to advise Congress on whether then-emerging technologies necessitated a change in copyright laws.¹³⁸ CONTU concluded that computers were, at least at that time, merely tools for facilitating human creativity.¹³⁹ According to this approach,

12/10/ai_intellectual_property_rights_how_artificial_intelligence_might_monetize.html [https://perma.cc/Z2D4-YN7K].

133. See Yanisky-Ravid & Liu, *supra* note 38, at 6, 18 (discussing inventions being produced by AI systems).

134. See David C. Vladeck, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 WASH. L. REV. 117, 120-29 (2014); see also O’Brien v. Intuitive Surgical, Inc., No. 10 C 3005, 2011 WL 3040479, at *1, *3 (N.D. Ill. July 25, 2011) (granting summary judgment to surgical robot’s manufacturer).

135. See Copyright, Designs & Patents Act 1988, c. 48, §§ 3, 9 (Eng.).

136. *Id.* § 9(1).

137. *Id.* § 9(3).

138. U.S., FINAL REPORT OF THE NATIONAL COMMISSION ON NEW TECHNOLOGICAL USES OF COPYRIGHTED WORKS 3, 4 (1978), <http://www.digital-law-online.info/CONTU/PDF/index.html> [https://perma.cc/A5GC-446C] [hereinafter CONTU FINAL REPORT].

139. See *id.* at 45.

Generating Rembrandt

691

The computer, like [the] camera or [] typewriter, is an inert instrument, capable of functioning only when activated either directly or indirectly by a human . . . [and] affects the copyright status of a resultant work no more than the employment of a still or motion-picture camera, a tape recorder, or typewriter.¹⁴⁰

Entities, such as employers and firms, are thus entitled to copyright ownership as the transferees of those programmers.¹⁴¹

This Article criticizes this traditional approach and calls on policymakers to revisit copyright laws in light of already-existing advanced technology and the latest developments in AI systems.¹⁴² I argue that, inevitably, current copyright law will not be able to cope with AI systems' productivity and creativity.¹⁴³ One reason is that too many stakeholders are involved in the process of creating the AI system itself, with no one acting as the main contributor.¹⁴⁴ This point of view holds the contributors involved in the process as owners of the AI system, and thus the ones responsible for works generated by the AI system.¹⁴⁵

1. *Who Could the Owner Be?*

The candidates for ownership of, and subsequent accountability for, AI works vary from one case to another.¹⁴⁶ However, entitlement to these rights depends on each candidate's direct or indirect contributions to the AI system.¹⁴⁷ I claim that due to the multi-player model, most of the time, the candidates who are involved in the development and manufacture of the AI system do not meet the threshold of authorship.¹⁴⁸ The programming and algorithms used by robots and AI systems may be the work of many hands and may employ generative technologies that allow innovation at multiple

140. *Id.* at 44-45.

141. *Copyright Ownership: Who Owns What?*, *supra* note 10.

142. *See* Fischer, *supra* note 124 (noting that non-human systems will create copyrightable works).

143. *See id.*

144. *See* Yanisky-Ravid & Liu, *supra* note 38, at 20 (suggesting that multiple stakeholders in inventions created by AI systems disrupts the traditional patent process because there is no single inventor).

145. *See id.* (discussing inventions being produced by AI systems).

146. *See id.* (discussing ownership in the context of responsibility for infringement).

147. *See id.* In the case of *The Next Rembrandt*, one entity included all the players.

148. *See id.* (describing the multi-player model in regard to AI systems generating inventions).

layers.¹⁴⁹ These features of robotics and AI enhance unpredictability and complicate causal responsibility for what robots and AI systems do.¹⁵⁰ In addition to the AI system software programmers, there are (too) many players and stakeholders that contribute to the process of creating, designing, developing and producing the AI systems themselves, but not the product autonomously produced by the AI systems. Among others are the data suppliers, trainers, feedback suppliers, holders of the AI system, system operators, employers or investors, the public, and the government. The large number of players significantly weakens each player's individual contribution and thus the bond between the software programmers and the products produced by the AI systems. There are many options for who should own the works created by AI systems and, indeed, one role may overlap with another. The following discussion will focus on some of these players.

First, there are the programmers of the AI system. Second, there are the trainers or the data providers, who may be among the most important figures shaping the final functions of the AI systems. Third, there are the feedback providers, or individuals whose task is to provide the AI system with a signal that allows it to distinguish right from wrong and sometimes to select the best result from many random, meaningless results.¹⁵¹ Fourth, there is the AI system's owner, whether that system is hardware or software. The owner might be the corporation, as the owner of the hardware (robot) or the software, or it might be the buyer of the AI systems (or robots). Fifth, there is the operator of the AI system, or the person who activates the system and enables its creation (although, it should be noted, some advanced AI systems can operate by themselves without a human operator).¹⁵² If one applies a practical approach, the operator could also be the manufacturer.¹⁵³ Sixth, there is the buyer of the

149. See Balkin, *supra* note 100, at 53 (noting that AI has innovation at multiple layers).

150. See *id.* (discussing causal responsibility of AI based on multiple hands working on programming and algorithms).

151. See Abbott, *supra* note 93, at 1082 (arguing "a computer's owner should be the default assignee of any invention, both because this is most consistent with the rules governing ownership of property, and because it would most incentivize innovation"); Weeks, *supra* note 30, at 93.

152. See Samuelson, *supra* note 129, at 1205 (discussing the role of the programmer and the programmer's claims to ownership).

153. See generally RICHARD T. WATSON, INFORMATION SYSTEMS (2012) (explaining the roles of manufacturers in AI systems).

Generating Rembrandt

693

product.¹⁵⁴ Seventh, the government or governmental entities could be entitled to ownership of products as a default or as a representative of the public. Eighth, the public could also be one of the candidates for ownership in cases of public domain policy.¹⁵⁵ Furthermore, different paradigms of ownership can exist regarding the suggested owners of works created by AI systems. In regard to these options, ownership could be sole ownership by one player or co-ownership by multiple stakeholders.

I argue that none of the players are entitled to ownership of the works generated by AI systems nor are they accountable for these works. Because of the features of AI systems—creative, autonomous, unpredictable, and evolving—none of the players can directly claim ownership and accountability of the works generated by AI systems. Furthermore, there are too many players involved in the process, and none of the players are the main contributor to the creation of the work. For example, although data and feedback providers are crucial to the process, they cannot be considered as owners because they are not authors. Thus, only one figure—the programmer—remains as a candidate for ownership and accountability.¹⁵⁶

2. *Distinguishing Between the Rights over Artificial Intelligence Software; the Rights of Works Produced by Automated AI Systems; and the Rights of Programmers*

For traditional artworks, the creators (or, in some cases, their employers or main contractors) are entitled to copyright over the artworks they produce, subject to several conditions.¹⁵⁷ As discussed above, developing the next generation of creative AI systems involves many participants, including software programmers and the

154. See Samuelson, *supra* note 129, at 1207-08.

155. See Muzdalifah Faried Bakry & Zhilang He, *Autonomous Creation – Creation by Robots: Who Owns the IP Rights?*, MAASTRICHT U. (Mar. 5, 2015), <https://law.maastrichtuniversity.nl/ipkm/autonomous-creation-creation-by-robots-who-owns-the-ip-rights> [<https://perma.cc/2YCC-RPER>] (arguing that artificial intelligence belongs in the public domain); Natasha Lomas, *We Need To Talk About AI and Access to Publicly Funded Data-Sets*, TECHCRUNCH (July 9, 2016), <https://techcrunch.com/2016/07/09/we-need-to-talk-about-ai-and-access-to-publicly-funded-data-sets/> [<https://perma.cc/G7KZ-97CN>] (explaining public domain data on Google).

156. See Samuelson, *supra* note 129, at 1205 (discussing the role of the programmer and the programmer's claims to ownership).

157. See *Copyright Ownership: Who Owns What?*, *supra* note 10.

companies for which they are working or those who commissioned the algorithm that generated the work, trainers that provide the data, and many other contributors.¹⁵⁸ The work itself, however, might be created digitally by an AI system embedded in a computer. I argue that the programmers of the software may be entitled to the copyright of the program, but may not necessarily have the rights for future products created by the AI system. I support this claim both conceptually and legally.

Conceptually, I argue that AI systems reflect a discipline focused on three inter-related components that are similar to the “human” traits of intelligence. First, unlike traditional software, the similarities and interconnections that the AI systems identify or find, process, remember, use, and implement may, in many cases, be unknown to the programmer. Second, in contrast to fixed and framed software, the AI system evolves and develops as a result of new input and new results. Third, the AI system’s works are significantly unpredictable because the system constantly and automatically evolves through its experiences.¹⁵⁹ In short, because of their intelligence components, AI systems are not only more accurate, of higher quality, and faster at processing details, but are also capable of creating unpredictable, original, and creative artworks and other products—all of which are unknown to their programmers. Therefore, these works created by AI systems could have been copyrightable under U.S. copyright laws.¹⁶⁰

Legally, the rights of an AI software program and the rights of artworks can be distinguished from one another. Software is usually protected not only by copyright laws, but also by the Constitution of the United States,¹⁶¹ which grants exclusive rights to “Authors and Inventors” in their respective “Writings and Discoveries.”¹⁶² However, the discourse about software ownership is distinct from the question of ownership of products (and services) produced by AI systems. One question that remains is whether the works produced by AI systems should or could be entitled to copyright protection. Can AI-generated works be regarded as proper “works of authorship” pursuant to § 102 of the Copyright Act by virtue of AI’s

158. See *supra* Part I (listing the AI participants).

159. See, e.g., Bridy, *supra* note 129, at 20 (explaining requirements for copyrighting); see also *supra* note 66.

160. See, e.g., Bridy, *supra* note 129, at 20 (explaining requirements for copyrighting).

161. See U.S. CONST. art. I, § 8, cl. 8.

162. See *id.*

Generating Rembrandt

695

sufficient nexus to human creativity?¹⁶³ Should this protection (if it does exist) also be applied to inventions produced by AI systems?¹⁶⁴ On the one hand, I do not challenge the programmers' entitlement to copyright ownership in the software they develop, but on the other hand, I argue that the entitlement to the software does not automatically result in ownership over the products created by AI systems.¹⁶⁵ I further conclude that the software programmers are not the owners of the works produced by AI systems, just as the owner of a brush or a camera does not hold the rights over the painting or the photo produced by those objects.

The distinction between programming the AI software itself and authoring the works the automated AI machine creates can be better understood by thinking about a piano and the author of the melodies created by using the piano. Imagine a melody that is created by Z playing a piano that was programmed and designed by A, manufactured by B, and owned by C. Is the piano (or the ownership of the piano) as the musical instrument, serving as the platform for the creation, relevant to the question of ownership of the melody?¹⁶⁶ I argue that neither the person who produced (or invented) the piano nor the factory that produced it are the owners of the melody created by a third entity (whether a human or an AI system).

Another relevant example would be the well-known selfie taken by a monkey with someone else's camera.¹⁶⁷ In this example, a monkey on the Indonesian island of Sulawesi took a photograph using a camera owned by David Slater, a nature photographer.¹⁶⁸ But

163. See 17 U.S.C. § 102 (1990).

164. See § 107 (under Copyright's "fair use" doctrine, others can reproduce the copyrighted inventions for "criticism, comment, news reporting, teaching . . . scholarship, or research"); Thomas Caswell & Kimberly Van Amburg, *Copyright Protection on the Internet*, in E-COPYRIGHT LAW HANDBOOK 7-1, 7-8 (Laura Lee Stapleton ed., 2003) (arguing that all who independently create inventions might be entitled to patent rights in order to protect it); Donald S. Chisum, *The Patentability of Algorithms*, 47 U. PITT. L. REV. 959, 1015-16 (1986).

165. See RICHARD STIM, *GETTING PERMISSION* 194 (6th ed. 2016); see also *Copyright Ownership: Who Owns What?*, *supra* note 10.

166. But see Pamela Samuelson, *Benson Revisited: The Case Against Patent Protection for Algorithms and Other Computer Program-Related Inventions*, 39 EMORY L.J. 1025, 1148 (1990) (arguing that the role of the software programmer is crucial).

167. Camila Domonoske, *Monkey Can't Own Copyright to His Selfie, Federal Judge Says*, NPR (Jan. 7, 2016), <http://www.npr.org/sections/thetwo-way/2016/01/07/462245189/federal-judge-says-monkey-cant-own-copyright-to-his-selfie> [<https://perma.cc/5N7J-YKZ5>].

168. See *id.*

Slater didn't trip the shutter—the monkey did.¹⁶⁹ The People for the Ethical Treatment of Animals (PETA) filed a lawsuit on behalf of the monkey, arguing that Naruto, the monkey, owns the copyright, which PETA offered to administer on the monkey's behalf.¹⁷⁰ Since the dispute began, “[t]he U.S. Copyright Office . . . has specifically listed a photograph taken by a monkey as an example of an item that cannot be copyrighted. Slater, meanwhile, has a British copyright for the photo, which he argues should be honored worldwide.”¹⁷¹ He has asked the U.S. court to dismiss PETA's claim.¹⁷² “Imagining a monkey as the copyright ‘author’ in Title 17 of the United States Code is a farcical journey Dr. Seuss might have written,” according to Slater's lawyer.¹⁷³

I argue that the producer or the seller of the instrument that served as the platform for producing new works (i.e., the camera, piano, or paintbrush)—like the software programmers or the companies in charge of producing the platform—are unsuitable candidates for being the creators or stakeholders of the works generated by the platform.¹⁷⁴ The owner of the work is the entity that generated the work. I argue that the rights to the AI systems' algorithms, which can be owned by the human creator, are distinct from the rights to the artworks the systems produce.

The code itself will have copyright protection. One could make the claim that the output generated from the computer program is a derivative work product of the underlying copyrighted program, which may also provide copyright protection to whomever holds a copyright in the algorithm. Thus, the holder of the copyright for the algorithm would hold the copyright for the output too.¹⁷⁵ However, in 1973, the Supreme Court interpreted the authorship requirement of the Copyright Act to include “any physical rendering of the fruits of

169. *See id.*

170. *See id.*

171. *See id.*

172. *Id.* (describing how, according to Slater's lawyer, “[t]he only pertinent fact in this case is that Plaintiff is a monkey suing for copyright infringement”).

173. *Id.*

174. *See* *Naruto v. Slater*, No. 15-cv-04324-WHO, 2016 WL 362231, at *1 (N.D. Cal. Jan. 28, 2016); Sarah Jeong, *Judge Gives Monkey Second Chance to Sue for Copyright Infringement*, MOTHERBOARD (Feb. 1, 2016, 3:40 PM), <http://motherboard.vice.com/read/judge-gives-monkey-second-chance-to-sue-for-copyright-infringement> [<https://perma.cc/AG8N-DPMY>] (discussing the judge's decision to give PETA leave to amend the complaint and try again to get damages from Slater).

175. *See* 17 U.S.C. § 101 (2012).

Generating Rembrandt

697

*creative intellectual or aesthetic labor.*¹⁷⁶ The Court concluded that, in most cases, a computer requires a significant amount of input from a human user in order to generate artistic output.¹⁷⁷

I argue that when the computer produces most of the output independently and creatively, it is less likely that the output might be considered to be the original source of the work and not as derivative work. I do *not* oppose the programmer's entitlement to ownership of the AI system itself. However, I do contest the human behind the machine's point of view and the idea that this entitlement automatically results in the programmer owning the products and processes created by the AI system.¹⁷⁸ I claim that my conclusion influences other cases beyond the intellectual property arena.¹⁷⁹ This brings me to another scenario targeting the AI system itself as being responsible for its own works.

3. *Other Possible Accountable Entities*

In other legal regimes, scholars have suggested strict liability as a solution for addressing the damages caused by AI systems, without blaming either the AI system or its programmers.¹⁸⁰ Strict liability is often employed when it would be too complicated to prove guilt, negligence, or a causal link between the defendant's failure to exercise due care and the damages that occurred.¹⁸¹ I argue that, due to the autonomous, creative, and unpredictable nature of AI systems, using the traditional strict liability rule on individuals would be unjust and inefficient.

Another option is to target the government or governmental body as being accountable.¹⁸² In some fields, such as international

176. *Goldstein v. California*, 412 U.S. 546, 561 (1973) (emphasis added).

177. *See id.*

178. *See infra* Section II.C (discussing the human behind the machine point of view and why this Article is critical of it).

179. *See infra* Subsection II.C.3 (discussing accountability implications of the idea that an AI creator does not necessarily own the AI's output).

180. *See* Vladeck, *supra* note 134, at 146.

181. *See, e.g., id.* ("My proposal is to construct a system of strict liability, completely uncoupled from notions of fault for this select group of cases. A strict liability regime cannot be based here on the argument that the vehicles are 'ultra-hazardous' or 'unreasonably risky' for the simple reason that diver-less vehicles are likely to be far less hazardous or risky than the products they replace."). *See also* Crotoft, *supra* note 11, at 1394-95 (arguing that autonomous weapons are designed to kill and their independent actions break the chain of causality, thereby making the strict liability rule applicable).

182. *See* Scherer, *supra* note 48, at 394.

law and autonomous weapons, the state is in the best position, at a practical level, to ensure compliance with the law (e.g., that autonomous weapons systems are designed and employed in compliance with international law).¹⁸³ States also have deep enough pockets to pay damages to the victims, in addition to being involved in developing, purchasing and using AI systems.¹⁸⁴ According to the proposed model, states as employers or users would bear responsibility for AI systems not because they are states per se, but rather for the reasons mentioned above, due to their status as users.¹⁸⁵ I argue that, at the national level, unlike the international level, responsibility could be forced. There may also be third party accountability.¹⁸⁶ In these solutions, accountability is not necessarily connected to ownership because the works generated by AI systems can be public domain, and copyrights laws may thus not be applicable at all.¹⁸⁷

I think that, under the copyright regime, these solutions do not efficiently serve the goal of imposing accountability on the player who should—along with enjoying the benefits of using AI systems—also take responsibility for such systems. I have discussed two alternative points of view.¹⁸⁸ First, the AI systems themselves could be the owners and the ones responsible for their works.¹⁸⁹ Second, the humans behind the machine (i.e., those involved in the process of developing the AI systems) could be the owners and the ones responsible for works generated by AI systems.¹⁹⁰ Since neither of these perspectives seems applicable and justified to the questions of ownership and accountability, I now turn to addressing these issues under a theoretical justification framework.

183. See Crotoof, *supra* note 11, at 1390.

184. See, e.g., Scherer, *supra* note 48, at 357, 394 (“This article will advance the discussion regarding the feasibility and pitfalls of government regulation of AI by examining these issues and explaining why there are, nevertheless, some potential paths to effective AI regulation.”). See also Crotoof, *supra* note 11, at 1389-93 (arguing that states are reluctant to take responsibility regarding autonomous weapons).

185. See Crotoof, *supra* note 11, at 1390.

186. See, e.g., Vladeck, *supra* note 134, at 148.

187. Yanisky-Ravid & Liu, *supra* note 38, at 18-21 (suggesting that inventions produced by AI systems will not be protected by the patent law).

188. See *supra* Part II.

189. See *supra* Section II.B.

190. See *supra* Section II.C.

Generating Rembrandt

699

III. THEORETICAL JUSTIFICATIONS

Many intuitively feel that AI systems, sophisticated robots, and machines should not be able to have rights and duties; nor should they hold copyrights. This intuition has its roots in strong theoretical and legal arguments.¹⁹¹ The following discussion will explain the difficulties of seeing AI systems as totally independent from human control. The discourse concerning the justifications for intellectual property focuses on three main substantive theories: law and economics, which examines intellectual property rules according to their cumulative efficiency and ability to promote total welfare; personality theory, which focuses on the personality of the creators; and Lockean labor theory, which justifies the property interest as the fruits of the creator's labor.¹⁹² Today, U.S. intellectual property law is based primarily on the law and economics utilitarianism approach¹⁹³ and, in part, John Locke's theory of labor.¹⁹⁴ By contrast, the civil law approach to copyright protection justifies property rights by the importance of the creators' personality in the works (personality approach), as well as by the ownership of the fruits stemming from the person's body and soul (Locke's approach or labor approach).¹⁹⁵

191. See Shlomit Yanisky-Ravid, *The Hidden Though Flourishing Justification of Intellectual Property Laws: Distributive Justice, National Versus International Approaches*, 21 LEWIS & CLARK L.R. 1, 8-9 (2017).

192. See *id.* at 4-9 (describing the three major approaches to theoretical justifications to intellectual property laws and arguing that distributive justice theory, although discussed by some scholars, is wrongfully considered to be neither a substantial nor a major justification of intellectual property; it is rather seen as an exception or postscript to the mainstream theoretical justifications). See also William Fisher, *Theories of Intellectual Property*, in NEW ESSAYS IN THE LEGAL AND POLITICAL THEORY OF PROPERTY 168, 169-75 (Stephen R. Munzer ed., 2001) (describing various theories underlying intellectual property); Justin Hughes, *The Philosophy of Intellectual Property*, 77 GEO. L.J. 287, 288-89 (1988) (discussing the different justifications to intellectual property laws).

193. DONALD S. CHISUM ET AL., PRINCIPLES OF PATENT LAW 50 (3d ed. 2004) (“[T]he predominant justification for American intellectual property law has been . . . utilitarianism.”).

194. Peter M. Kohlhepp, *When the Invention Is an Inventor: Revitalizing Patentable Subject Matter to Exclude Unpredictable Processes*, 93 MINN. L. REV. 779, 781-82 (2008).

195. See Jeanne C. Fromer, *Expressive Incentives in Intellectual Property*, 98 VA. L. REV. 1745, 1746 (2012) (discussing the personality and labor approach to intellectual property); Yanisky-Ravid, *supra* note 34, at 118.

A. Law & Economics

The U.S. system of copyright laws was established to protect original authors and creators by giving them exclusive rights and control over the works they generate.¹⁹⁶ The U.S. Constitution grants Congress the power “To promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries.”¹⁹⁷ The main justification for establishing a copyright regime—giving stakeholders property rights, which are broader than the rights established by the contract regime—is based on the theory of law and economics. In short, providing stakeholders property rights promotes the creation of useful art.¹⁹⁸ This, in turn, motivates the creators (or their transferees) to create, expose, develop, and distribute their works, enriching the total welfare of the public.¹⁹⁹ The Copyright Clause, by securing exclusive rights to authors and inventors, aims to “promote the . . . science and the useful Arts.”²⁰⁰

This Section will focus on copyright law’s purpose of promoting the creation of artistic works by establishing an incentive structure through which authors are given exclusive control over the copyright works.²⁰¹ Often, however, as a result of a special contract or relationship with the author or creator, other entities are entitled to the copyrights as direct transferees of the actual human creator.²⁰²

Unlike humans, AI systems do not need incentives to create artworks.²⁰³ It’s true that programmers need to be incentivized to

196. U.S. CONST. art I, § 8, cl. 8.

197. *Id.*

198. *Id.*

199. Richard A. Posner, *Intellectual Property: The Law and Economics Approach*, 19 J. ECON. PERSP. 57, 60 (2005) (explaining how the society made an agreement with the authors to grant them exclusive rights for limited duration and then the rights become public domain).

200. U.S. CONST. art. I, § 8, cl. 8.

201. See Posner, *supra* note 199, at 57; see also 17 U.S.C. § 106 (2017) (“Exclusive rights in copyrighted works.”); William M. Landes & Richard A. Posner, *An Economic Analysis of Copyright Law*, 18 J. LEGAL STUD. 325, 326 (1989).

202. SHLOMIT YANISKY-RAVID, *INTELLECTUAL PROPERTY IN THE WORKPLACE: THEORETICAL AND COMPARATIVE PERSPECTIVES* (2013) (explaining the incentives to create as being part of the law and economics justification as well as other justifications for intellectual property); Pamela Samuelson, *Mapping the Digital Public Domain: Threats and Opportunities*, 66 L. & CONTEMP. PROBS. 147, 156 (2003).

203. See *supra* note 32 and accompanying text.

Generating Rembrandt

701

create and develop advanced, automated AI systems, but programmers, or the entities for which they are working, do hold copyrights over the software.²⁰⁴ Once we understand the nature of incentives, we understand that they are nevertheless needed to (1) promote the development of AI systems' programming and (2) encourage entities to control the functions of AI systems and to take responsibility for their outcomes. In these cases, ownership might be the most efficient tool for gaining this incentive.²⁰⁵

However, we do not need to incentivize robots or AI systems to function. Incentivizing AI systems to generate works they are already internally programmed to create is pointless. My argument is rooted in understanding that automated AI systems not only evolve independently after the program has been completed, but also evolve in ways that are unpredictable, even to the human programmers who created them. This conclusion is further drawn from the fact that the connection and similarities that AI systems draw are neither made nor known to the programmers.²⁰⁶ We can compare this system to human perception via the human brain. The programmers implemented or created the neurons and synapses, but not the electronic messages that will be created in the future and their products.²⁰⁷ The programmers thus create the systems, but cannot predict the works themselves.²⁰⁸ Consequently, the creativity of an AI system is not a result of the creativity of the programmers; at the very least, the causal relationship is not close enough to justify ownership (as a tool to incentivize a specific function) in the new works generated by AI machines.²⁰⁹ The human programmer is only ancillary to the creation of the artworks.

204. See Yanisky-Ravid & Liu, *supra* note 38, at 15.

205. Garrett Hardin, *The Tragedy of the Commons*, SCIENCE, 1243, 1243-48 (1968) (arguing that ownership is efficient to retain property).

206. See Yanisky-Ravid & Liu, *supra* note 38, at 2.

207. Julien Vitay, Helge Ü. Dinkelbach & Fred H. Hamker, *ANNarchy: A Code Generation Approach to Neural Simulations on Parallel Hardware*, FRONTIERS NEUROINFORMATICS, July 31, 2015, at 1 (discussing a notable exception, the Brian simulator, "which allows the user to completely define the neuron and synapse models using a simple mathematical description of the corresponding equations [and] uses a code generation approach to transform these descriptions into executable code, [which in turn] allow[s] the user to implement any kind of neuron or synapse model").

208. JAMES GLEICK, WHAT JUST HAPPENED: A CHRONICLE FROM THE INFORMATION FRONTIER 19 (2002).

209. See Yanisky-Ravid & Liu, *supra* note 38, at 18-19.

In addition, and as mentioned above, programmers are already incentivized to make creative AI systems because they receive copyright protection for the program itself.²¹⁰ Furthermore, because copyright protection does not exist in a vacuum, it must be balanced against competing rights. It is important that the legal regime incentivizes the right people and entities, and ultimately promotes behavior that will increase total welfare.²¹¹ The legal regime has succeeded if programmers who create AI systems are incentivized to do so either through intellectual property protection or patent protection for the machine, copyright protection for its computer code, or both. But if we understand that these legal tools incentivize the AI system or the programmers to create works of authorship, when they are not in fact doing so, the system is failing because it is inefficient. It should be obvious that machines need no incentive to work. In other words, assuming that machines capable of creating unique art already exist, in all likelihood there would be no need to incentivize the creation of these works. Providing AI systems with wires, electronic devices, Internet connection, and materials should be enough.

If, as the law and economics approach contends, copyright is meant to be an incentive structure, and machines do not need to be incentivized to create, then copyrighting the machines' works provides no benefit but does hamper the public's ability to enjoy the work.²¹² Thus, giving AI systems rights to the works they create would seemingly operate to take them out of the sphere of copyright altogether.²¹³ Indeed, the public's or the end-users' interest in appreciating and enjoying works of art should be balanced against the private interest in maintaining exclusive, monopolistic control.²¹⁴ Since human creators need to be incentivized to create, copyright used to be the optimal state of affairs for both parties because, without it, much fewer works of art would be created for the public

210. See *id.* at 15.

211. Tiina Kautio et al., *Assessing the Operation of Copyright and Related Rights Systems*, CUPORE (2016), <http://www.cupore.fi/en/research/previous-researches/assessing-the-operation-of-copyright-and-related-rights-systems-142507-14122016> [<https://perma.cc/4BM3-SNCA>].

212. See Yanisky-Ravid & Velez-Hernandez, *supra* note 32, at 7; see also Yanisky-Ravid & Liu, *supra* note 38, at 29-30 (arguing that patent laws are not applicable in the 3A era of AI).

213. See Yanisky-Ravid & Liu, *supra* note 38, at 8.

214. See Julius Cohen, *The Anti-Trust Acts and "Monopolistic Competition": A Case Study*, 24 CORNELL L. REV. 80 (1938).

Generating Rembrandt

703

to enjoy.²¹⁵ But in the non-hypothetical future in which machines can create pleasing works of art without limits, I argue that the existing balance would be thrown off. In the case of AI systems, I argue, there wouldn't be any risk of a lack of artistic creation even if copyright law did not exist to protect such creations.²¹⁶ Such a reality could, furthermore, pose an existential threat to the entire copyright regime.²¹⁷

Assuming that many people consume works of authorship for their artistic merit, I argue it is likely that machine-produced works could not serve as a perfect replacement for human-authored works. Instead, the market for human-authored works of art would coexist with a market for works "authored" by machines.²¹⁸ Since human artists would still need to be compensated, copyright law would persist, at least until machines capable of imparting deeper meaning to their work were created (if such a thing is indeed possible).²¹⁹ In addition to being more likely in the near future, this model is perhaps more palatable to policymakers and the general public.

Thus, denying copyright protection for works of authorship created by machines is unlikely to greatly change the existing system. However, as the world becomes more electronically based and cyber-focused (a trend we can already observe), it won't take long until machines, using AI systems, can copy any artistic work precisely (including the signature).²²⁰ This will ultimately destroy the

215. See Jared Green, *Why Public Art Is Important*, DIRT (Oct. 15, 2012), <https://dirt.asla.org/2012/10/15/why-public-art-is-important> [<https://perma.cc/T7NX-E8PA>].

216. See *Moral Rights*, ARTS L. CTR. AUSTL., <http://www.artslaw.com.au/info-sheets/info-sheet/moral-rights> [<https://perma.cc/8YNQ-67PA>] (last visited Jan. 15, 2018). Moral rights protect the personal relationship between a creator and his or her work even if the creator no longer owns the work or the copyright in the work. Moral rights concern the creator's right to be properly attributed or credited and the protection of his or her work from derogatory treatment. See *id.*

217. See Yanisky-Ravid & Velez-Hernandez, *supra* note 32, at 19. See generally Tang, *supra* note 34 (explaining how the involvement of digital tools in creation leads to seeing moral right as trademark).

218. See Samuelson, *supra* note 129; see also Samuelson, *supra* note 166, at 1148 (arguing that the role of the software programmer is crucial).

219. See generally *Artificial Intelligence – Overview*, TUTORIALSPPOINT, https://www.tutorialspoint.com/artificial_intelligence/pdf/artificial_intelligence_overview.pdf [<https://perma.cc/5M4P-WVZB>] (last visited Jan. 15, 2018).

220. See e.g., Shlomit Yanisky-Ravid & Kenneth S. Kwan, *3D Printing the Road Ahead: The Digitization of Products when Public Safety Meets Intellectual Property Rights—A New Model*, 38 CARDOZO L. REV. 921, 923-24 (2017).

incentive to produce these works of art, which, in turn, will eventually destroy the copyright regime.²²¹

One possibility is that the AI systems might require electronic licenses, drawn up by electronic agreements,²²² to use their products, as well as electronic contracts creating electronic sanctions for breaching the license (e.g., electronically terminating the infringing works).²²³ However, these methods would not need copyright laws, as the theoretical rights and their enforcement would no longer use the traditional court system.²²⁴ Although AI systems might be able to detect infringements easier and in more efficient ways, implementing copyright laws for the purpose of excluding other entities is not the right solution. Doing so would most likely lead to the loss of control and lack of accountability and responsibility that humans have over property and intellectual property rights.²²⁵

The thought of machines taking over and nullifying copyright law is not just far-fetched; it would also require a tremendous, uncomfortable shift in the legal landscape. After confronting the challenges posed by advanced technology and AI systems that autonomously generate works, it would be a stretch—even in the existing case of a sophisticated neural network AI capable of learning and creating independently—to imagine an AI system that could understand and use the copyright regime as its incentive. Furthermore, it seems non-feasible that AI systems will be capable in the near future of suing in court for ownership rights.²²⁶ I contend that, while preparing and formulating future laws, although theoretically and digitally feasible, it is not likely that AI systems will acquire ownership and sell or give licenses to use their products in the near future. I further claim that even when AI systems will be qualified to possess their own rights and duties, a more theoretically

221. *See id.* at 927.

222. *See* Scholz, *supra* note 17, at 102.

223. *See id.* at 110.

224. *See id.* at 120-21.

225. Yanisky-Ravid & Kwan, *supra* note 220, at 924 (discussing the threat and hazards of 3D printings).

226. The decision earlier this month in the case of *Halo v. Pulse* will give owners of U.S. patents a greater likelihood of being awarded enhanced damages. *See* Frederic Henschel & Kevin M. Littman, *U.S. Supreme Court Strengthens Patents (for a Change)*, SCIENCE BUS. (June 23, 2016), [http://sciencebusiness.net/news/79833/US-Supreme-Court-strengthens-patents-\(for-a-change\)](http://sciencebusiness.net/news/79833/US-Supreme-Court-strengthens-patents-(for-a-change)) [<https://perma.cc/UCX3-3WUP>] (arguing this will raise the value of patents and increase the incentive to sue for infringement).

Generating Rembrandt

705

justified solution will be to legally impose these rights and duties on other parties as the users.

Instead, I would like to suggest an alternative model that, on the one hand, acknowledges and reflects the perception of the 3A era of automated, autonomous, and advanced AI systems, and, on the other hand, imposes control and accountability on traditional legal entities. *This* model would consider AI systems as employees (or contractors) that work for the humans or firms that legally operate them. *This* model is similar to the notion of an “employed creator” under the WMFH doctrine—i.e., an employee who creates new works in the scope of their employment.

The owner of a copyright has the exclusive right and may authorize others to reproduce the work, prepare derivative works based on the work, distribute copies of the work, or show the work publicly. Having those rights also means that the copyright holder has the right to stop others from infringing on those rights. The problem for a non-human, such as an AI system, is that it is unable to enforce those rights. Although it is theoretically feasible, a computer cannot sue another computer in court over the unauthorized copying of its work. Furthermore, a computer is incapable of transferring those rights to others who might be able to sue on its behalf. Even from a public policy perspective, the main purpose of granting copyright protection is to stimulate artistic creation by ensuring that nobody can steal the fruits of an artist’s labor, making it less risky to create original works of authorship. Since computers cannot be “encouraged” to create new works, the usual public policy justifications underlying copyright law are inapplicable.

Some would argue that the WMFH model isn’t any different from a film director and a cameraman taking particular shots. The cameraman is a creative person, but the director will hold the right to the shot. AI systems act similarly to the creative cameraman. In fact, in *Goldstein v. California*, the Supreme Court interpreted the authorship requirement to include “any physical rendering of the fruits of *creative intellectual or aesthetic labor*.”²²⁷ The Court reasoned that, in most cases, in order for a computer to generate any kind of artistic work, it would require significant input from an author or user.²²⁸ Another way to think about it is this: when an artist uses Adobe Illustrator to create a unique graphic design, nobody can deny that the designs were the product of the designer’s creative

227. 412 U.S. 546, 561 (1973) (emphasis added).

228. *See id.*

mind. However, creating a song by pressing a button on a random number music generator isn't going to receive copyright protection on the resulting musical composition. But if the user provides some input that affects the song being generated, such as choosing the instruments, deciding on the key or tempo, or choosing a musical style for the composition, then the final musical composition may be the result of creative input and therefore copyrightable.

The law and economics theory, discussed above, is the dominant justification for copyright protection in Anglo-American law.²²⁹ However, in continental Europe, where copyright protection originated with an eye towards protecting great, independent artists, a different approach prevails, as addressed in the next Section.

B. Personality and Labor Theoretical Justifications

In civil law jurisdictions, the dominant justifications for copyright are the personality and labor Lockean theories.²³⁰ The personality theory posits that copyright protection is a right that accrues to the author in possession, reflection, and development of his personality on the assets.²³¹ It recognizes and appreciates the author's accomplishments and the element of his or her personality and individuality that the work contains, rather than simply an incentive to create more.²³² A related justification is the labor theory, which stipulates that copyright protection exists due to the hard work and dedication that authors contribute to their works.²³³ Just as AI does not need to be incentivized, AI systems do not have any need for recognition of the works reflecting their personality.²³⁴ Nevertheless, I argue that copyright protection could still accrue to the creators of such machines.

Developing AI systems capable of creating works of authorship is a great accomplishment. Therefore, it may make sense to grant

229. See Posner, *supra* note 199.

230. See MARTIN SENFTLEBEN, COPYRIGHT, LIMITATIONS, AND THE THREE-STEP TEST: AN ANALYSIS OF THE THREE-STEP TEST IN INTERNATIONAL AND EC COPYRIGHT LAW 8 (2004); Justin Hughes, *The Personality Interest of Artists and Inventors in Intellectual Property*, 16 CARDOZO ARTS & ENT. L.J. 81, 83 (1998).

231. See Hughes, *supra* note 230, at 83.

232. See Margaret Jane Radin, *Property and Personhood*, 34 STAN. L. REV. 957, 986 (1982) (arguing that the more personal one's property is, the more nonfungible and nontransferable it becomes); see also Yanisky-Ravid, *supra* note 191, at 9.

233. Yanisky-Ravid, *supra* note 191, at 4-5.

234. See *supra* Part III.

Generating Rembrandt

707

programmers the copyrights of works created by AI systems to recognize the magnitude of that accomplishment. Still, and even more strongly than the analysis of the incentive structure endemic to the law and economic theory, we have no other option than to recognize that when a creator is a machine, robot, or AI system, the personality theory and the labor theory are irrelevant. We therefore cannot justify the existence of copyright laws when they are applied to this new reality. Just as we do not need to incentivize programmers to create works of authorship in which they do not have any creative input, we do not need to recognize a programmer for an artistic accomplishment that is not his or her own. Therefore, there is little support for granting copyright protection to human programmers for the works of their AI systems under this theory either. However, when addressing the Work Made for Hire doctrine, we can rely partially on the labor approach to the investment of the firm in the works produced by the AI systems.²³⁵

IV. THE MODEL OF AI—WORK MADE FOR HIRE (WMFH)

One major motivation for the proposed model is to unveil the clandestine interests behind the phenomenon of AI systems. Following Professor Jack Balkin, who has explored the “laws of robotics” and the legal and policy principles that should govern how human beings use robots, algorithms, and AI systems,²³⁶ I claim that we should view AI systems as working for the users, and hence the users should bear accountability for the systems’ production, in addition to the benefits thereof. Balkin argues that there exists a false belief of a little person inside each robot or program who has either good or bad intentions.²³⁷ According to Balkin, the substitution effect refers to the multiple effects on social power and social relations that arise from the fact that robots, AI systems, and algorithms act as substitutes for human beings and operate as

235. See, e.g., Shlomit Yanisky-Ravid & Amy Mittelman, *Gender Biases in Cyberspace: A Two-Stage Model, the New Arena of Wikipedia and Other Websites*, 26 *FORDHAM INTELL. PROP. MEDIA & ENT. L.J.* 381, 391 (2016) (explaining that investment into information technology can help establish a more free and democratic reality); see also Shlomit Yanisky-Ravid, “*For a Mess of Pottage*”: *Incentivizing Creative Employees Toward Improved Competitiveness*, *CORNELL HUM. RTS. REV.* (2013), <http://www.cornellhrreview.org/wp-content/uploads/2013/05/Incentivizing-Creative-Shlomit-Yanisky-Ravid.pdf> [<https://perma.cc/88MX-T8WM>].

236. See Balkin, *supra* note 21, at 14.

237. See *id.* at 13-14.

special-purpose people.²³⁸ For Balkin, the most important issues in the laws of robotics require an understanding of how human beings exercise power over other human beings mediated through new technologies.²³⁹ The “three laws of robotics” should therefore be laws directed at human beings and human organizations, not at the robots or AI systems. According to Professor Balkin, those basic laws that regulate and control robots and AI systems include the following: (1) operators of robots, algorithms, and AI systems are information fiduciaries who have special duties of good faith and fair dealing toward their end-users, clients, and customers; (2) privately owned businesses who are not information fiduciaries nevertheless have duties toward the general public.²⁴⁰ I further argue that identifying the many players behind AI systems is the key factor for imposing accountability for the works generated by AI systems. Following Balkin’s argument, I propose a new model that might delegitimize the use of new technologies as a means for both public and private organizations to govern large populations. In order to unveil these hidden powers, I propose a model that sees AI systems as independent workers or employees of the users.

A. Rethinking the WMFH Legal Doctrine in the Case of AI Systems

The WMFH doctrine gives employers, or the individual commissioning the work, the copyright in works of authorship created by the employees or subcontractors.²⁴¹ The WMFH rule is thus an exception to the general principle of copyright ownership. Usually, the copyright becomes the property of the author once the creation meets the demands of the law.²⁴² However, if a work is made

238. *See id.* at 14.

239. *See id.* at 16.

240. *See id.* at 19-23 (arguing that those who use robots, algorithms, and AI systems have a public duty to avoid creating nuisances. Thus, for example, businesses may not leverage asymmetries of information, monitoring capacity, and computational power to externalize the costs of their activities onto the general public).

241. 17 U.S.C. §§ 101, 201 (2010); *Cnty. for Creative Non-Violence v. Reid*, 490 U.S. 730, 746 (1989); *Works Made For Hire*, U.S. COPYRIGHT OFF., <https://www.copyright.gov/circs/circ09.pdf> [<https://perma.cc/PTA4-X44R>] (last visited Jan. 15, 2018).

242. § 102 (“Subject Matter of Copyright (a) Copyright protection subsists, in accordance with this title, in original works of authorship fixed in any tangible medium of expression, now known or later developed, from which they can be perceived, reproduced, or otherwise communicated, either directly or with the aid of a machine or device. Works of authorship include the following categories: (1)

Generating Rembrandt

709

for hire, the employer or the one who commissioned the work would be considered the author, even if an employee or subcontractor actually created the work. The employer could be a firm, an organization, or an individual.²⁴³

Section 101 of the Copyright Act defines a “work made for hire” in two parts:

- (1) a work prepared by an employee within the scope of his or her employment; or
- (2) a work specially ordered or commissioned for use as a contribution to a collective work, as a part of a motion picture or other audiovisual work, as a translation, as a supplementary work, as a compilation, as an instructional text, as a test, as answer material for a test, or as an atlas, if the parties expressly agree in a written instrument signed by them that the work shall be considered a work made for hire.²⁴⁴

This section should be read together with Section 201 of the same Act:

(a) Initial Ownership.

Copyright in a work protected under this title vests initially in the author or authors of the work. The authors of a joint work are coowners of copyright in the work.

(b) Works Made for Hire.

In the case of a work made for hire, the employer or other person for whom the work was prepared is considered the author for purposes of this title, and, unless the parties have expressly agreed otherwise in a written instrument signed by them, owns all of the rights comprised in the copyright.²⁴⁵

The Supreme Court’s decision in *Community for Creative Non-Violence v. Reed* addressed the “work made for hire” definition.²⁴⁶

literary works; (2) musical works, including any accompanying words; (3) dramatic works, including any accompanying music; (4) pantomimes and choreographic works; (5) pictorial, graphic, and sculptural works; (6) motion pictures and other audiovisual works; (7) sound recordings; and (8) architectural works. (b) In no case does copyright protection for an original work of authorship extend to any idea, procedure, process, system, method of operation, concept, principle, or discovery, regardless of the form in which it is described, explained, illustrated, or embodied in such work.”)

243. See *Works Made for Hire*, *supra* note 241, at 2 (“If a work is made for hire, the employer or other person for whom the work was prepared is the initial owner of the copyright unless both parties involved have signed a written agreement to the contrary.”).

244. § 101.

245. § 201.

246. *Cnty. for Creative Non-Violence*, 490 U.S. at 737.

“The Court held that one must first ascertain whether a work was prepared by an employee or an independent contractor.”²⁴⁷ “If an employee created the work . . . the work will generally be considered a work made for hire.”²⁴⁸ In this context, however, the term *employee* differs from its common understanding.²⁴⁹ “For copyright purposes, ‘employee’ means an employee under the general common law of agency.”²⁵⁰ “An independent contractor,” on the other hand, “is someone who is not an employee under the general common law of agency.”²⁵¹ “If an independent contractor created the work, and the work was specially ordered or commissioned,” the second part of the WMFH definition applies.²⁵² “A work created by an independent contractor can be a work made for hire only if (a) it falls within one of the nine categories of works listed . . . above, and (b) there is a written agreement between the parties specifying that the work is a work made for hire.”²⁵³

To help determine who is an employee, the Court identified factors that establish an “employer–employee” relationship, as defined by agency law.²⁵⁴ The factors fall into three broad categories:

- (1) control by the employer over the work (i.e., the employer determines how the work is done, has the work done at the employer’s location, and provides the . . . means to create the work);
- (2) control by the employer over the employee (i.e., the employer controls the employee’s [time] in creating the work, has the right to have the employee perform other assignments . . . or has the right to hire the employee’s assistants); and
- (3) status and conduct of the employer (i.e., the employer is in business to produce such works [or] provides the employee with benefits).²⁵⁵

“These factors are not exhaustive[,] [and] [t]he Court left unclear which of these factors must be present in order to establish the employment relationship under the work-for-hire definition.”²⁵⁶

247. See *id.* at 731; *Works Made for Hire*, *supra* note 241, at 2.

248. See *Cnty. for Creative Non-Violence*, 490 U.S. at 732; see also *Works Made for Hire*, *supra* note 241, at 2.

249. See *Works Made for Hire*, *supra* note 241, at 2.

250. See *id.*

251. See *id.*

252. See *id.*

253. See *id.*

254. See *id.*

255. See *Cnty. for Creative Non-Violence v. Reid*, 490 U.S. 730, 751-52 (1989); see also *Works Made for Hire*, *supra* note 241, at 2.

256. See 17 U.S.C. §§ 101, 201 (2012); *Community for Creative Non-Violence*, 490 U.S. at 751-52.; *Works Made for Hire*, *supra* note 241, at 2.

Generating Rembrandt

711

Examples of works made in an [employer–employee] relationship include: [a] software program created by a staff programmer within the scope of his or her duties at a software firm[;] [a] newspaper article written by a staff journalist for publication in the newspaper that employs the journalist . . . [;] [a] musical arrangement written for a music company by a salaried arranger on the company’s staff[;] [and] [a] sound recording created by the salaried staff engineers of a record company.²⁵⁷

Why it is important to identify WMFH? There are important consequences that stem from the WMFH doctrine, including that the term and duration of copyright protection differ, there are no moral rights, and the termination provisions of the law do not apply.²⁵⁸

B. WMFH and Works Generated by AI Systems

This doctrine is an important and major exception to the general rule that copyright protection properly rests with the one or the many who actually created the work.²⁵⁹ It is therefore important for cases of AI systems generating works.²⁶⁰ The Copyright Act named the employer and main contractor as the authors of the work even though they have not actually created the work.²⁶¹ The policy rationale for this doctrine is to incentivize the employer or primary contractor at whose instance, direction, use, commercial purposes or risk the work is prepared, as well as to give them control over the commercial force regarding the work.²⁶² The idea and the outcome is that the employer or primary contractor, rather than the creator (who is an employee or sub-contractor), has the responsibility for and the accountability over the actions of the creator in regards to, inter alia,

257. See *Works Made for Hire*, *supra* note 241, at 2.

258. §§ 101, 106A, 302(a), 302(c), 304(a), 203(a). For example, WMFH copyright protection of a work made for hire is ninety-five years from the date of publication or 120 years from the date of creation, whichever expires first, whereas a work not made for hire is ordinarily protected by copyright for the life of the author plus seventy years. See *id.* § 302.

259. See § 201(a).

260. See Catherine L. Fisk, *Removing the ‘Fuel of Interests’ from the ‘Fire of Genius’*: *Law and the Employee-Inventor, 1830-1930*, 65 U. CHI. L. REV. 1127, 1131 (1998) (arguing that employees efficiently transfer their rights in future products to their employers through their employment contracts).

261. See § 201(b).

262. *Cmty. for Creative Non-Violence v. Reid*, 490 U.S. 730, 746 (1989); Fisk, *supra* note 260. For a critique of this approach, see Shlomit Yanisky-Ravid, *Rethinking Innovation and Productivity Within the Workplace Amidst Economic Uncertainty*, 24 FORDHAM INTELL. PROP., MEDIA & ENT. L.J. 143, 173-79 (2013) (rethinking innovation by incentivizing employees), and Yanisky-Ravid, *supra* note 235, at 3.

infringements of the law and harm caused by the work.²⁶³ This rule may be altered or changed by a contract among the relevant parties.

I claim that this doctrine seems to fit well conceptually with the problem of works created by AI systems. Although the AI system itself would be the proximate creator of the work, others, such as the user of the AI system at whose instance the work is ultimately created, will be entitled to ownership as well as accountability in regard to the works. But in the case of AI systems, who is the employer or main contractor? The answer may be complicated and may vary according to different circumstances. In many cases, it will be the user that operates and provides directions to the machine in the form of instructing it what to paint, write about, etc. The answer may also be the user that takes the financial risk of buying or hiring the machine and supplying it with energy and materials in the hope of producing a marketable final product. From a policy and practical standpoint, it makes sense to incentivize people or firms as well as other entities to use creative AI systems to create works of authorship because doing so will most efficiently promote the proliferation of the devices and the works they produce.²⁶⁴

The justification for giving the entitlement of ownership to economic entities is rooted in the incentive theory as well.²⁶⁵ This legally sanctioned monopoly allows the users to use, sell, or distribute the works more efficiently, as well as to be accountable for avoiding infringements and counterfeits.²⁶⁶ The latter is perhaps a better argument for giving copyright protection in the works of advanced, autonomous AI systems to their users. To avoid AI systems getting out of control, we have to legally nominate the most efficient entity to control them. The incentive for imposing property accountability on the users as employers or main contractors and seeing AI systems as employees or subcontractors is not just intuitive, it is also justified by theoretical and practical reasoning. The user can also be the owner of the AI system when the owner is the more efficient entity for controlling these works.

263. 35 U.S.C. §§ 101, 102 (2012); Fisk, *supra* note 260, at 1131.

264. See Yanisky-Ravid, *supra* note 235, at 3.

265. Kendra Cherry, *The Incentive Theory of Motivation: Are Actions Motivated by a Desire for Rewards*, VERY WELL (May 9, 2016), <https://www.verywell.com/the-incentive-theory-of-motivation-2795382> [<https://perma.cc/3JN7-P24S>].

266. Paul Belleflamme, *The Economics of Copyright Protection*, IPDIGIT (Oct. 2, 2013), <http://www.ipdigit.eu/2013/10/the-economics-of-copyright-protection> [<https://perma.cc/MGQ3-6QNW>].

Generating Rembrandt

713

This model also solves the inherent problem of multiple players being involved in the development of AI systems. The tragedy of multiple stakeholders is that they can block the development and commercial use of the AI system.²⁶⁷ Moreover, the model would encourage further investment in the AI industry and likely promote science and technology, thus promoting the goals of the Constitution and promoting total welfare. With respect to AI systems, the innovation provided by this model does not just grant rights and benefits, such as ownership of the products, it also imposes responsibility and thus assists in solving the problem of the lack of accountability for the outcomes of AI systems. This mechanism might also contribute to the responsibility and accountability for the use of AI systems in other regimes, such as criminal law and tort law. One could argue that these fields are based on a different justification and, therefore, are not influenced by the copyright regime. However, I claim that, because AI systems are copyrightable based on their software, it may be justified and useful to implement this model within the intellectual property realm as it intersects with other legal fields, such as tort and criminal law, that address the same challenges, including lack of accountability for damages generated by autonomous car accidents caused by AI systems.

Under this model, we see the AI systems as creative employees or subcontractors (just like humans) working for these entities. The model works for both firms and humans: The autonomous AI system, just like WMFH-employed creators, is the creative author of a work. When an AI system acts autonomously, it can be compared to an independent contractor and thus be shielded under WMFH doctrine.

C. The Legal Implications of the AI WMFH Model

Who owns the copyright in regard to the works generated by an AI system? Who is responsible for any damage the works may cause? Who would be the most efficient player in distributing and selling the works? Take, for example, The Next Rembrandt project. Unlike a traditional computer program, The Next Rembrandt project had teams of people working for several years to bring it to the public. What happens to those individuals? Do all of the people involved with the project have copyright ownership of its artworks?

267. Michael A. Heller, *The Tragedy of the Anticommons: Property in the Transition from Marx to Markets*, 111 HARV. L. REV. 621, 621-24 (1998).

Are they all, or perhaps only some of them, entitled to joint copyright ownership? Trying to determine the scope of ownership amongst the team members would be extremely difficult. In fact, this multi-stakeholders challenge was one of the practical and theoretical issues that led to the original WMFH doctrine.

On some level, the AI WMFH doctrine can solve this problem. It holds that the person or entity that orders or initiates the work is entitled to the copyright, instead of the authors themselves. Based on this theory, before the AI system was generated, the employer or the main contractor may be entitled to all of the rights. However, does this mean that the employer or the main contractor, under certain circumstances, is also entitled to the right over the paintings generated by the AI system? If this were the case, for example, the entity that operates The Next Rembrandt project, ING, would receive the full copyright over the paintings being generated by the system, as soon as certain legal requirements were met.²⁶⁸ Thus, it is possible that there is a copyright in The Next Rembrandt and that the copyright is held by ING. Copyright protection is only important if ING wants to enforce it, and applying the WMFH doctrine in a case like this would have some drawbacks.

D. The Drawbacks of Adopting the WMFH Model in Cases of AI Systems

Many questions arise in implementing the existing WMFH doctrine. Are the works generated by AI systems copyrightable in the first place? If these works are not copyrightable, can the employer hold copyright through the WMFH doctrine? What happens if the works generated by AI systems are not included in the nine-item list

268. The requirements being: (1) A written agreement signed by both parties (2) that specifically states that the work is a “work-made-for-hire” and (3) the work must be one of these nine types: a contribution to a collective work, part of a motion picture or other audiovisual work; a translation; a supplementary work; a compilation; an instructional text; a test; material for a test; or an atlas. Generally, in order for the WMFH doctrine to apply when many individuals are involved in producing a work, the entity entitled to copyright ownership must sign a contract with each team member attesting that each team member’s contribution is a work made for hire. The type of work must also be included in the list of products covered by the WMFH doctrine. An argument could be made that The Next Rembrandt might fall under the category of “compilation,” or perhaps a “contribution to a collective work.” Additionally, it is very likely that ING, with potential copyright claims to the work, had to affirmatively relinquish any claims prior to starting work on the project.

Generating Rembrandt

715

of the copyright law? What happens when autonomous AI systems create a work outside the scope of “employment”? What would be the legal outcome in another jurisdiction, such as France, where the creative employees retain the rights themselves? What would be the outcome when the AI system generates products or actions that are not copyrightable?

The Supreme Court has suggested that the WMFH doctrine is very limited in scope—namely, it applies only to instances where Congress has expressed a clear and explicit intent to override section 102.²⁶⁹ Therefore, implementing the doctrine would require new legislation with a broader scope of the matters and the rights involved. By comparison, denying copyright to works produced by advanced AI systems would probably require judicial clarification, as such a result is theoretically compatible with the current legal framework.

Furthermore, the AI context is less germane to the Work Made for Hire analysis than a corporation, like a publishing company or record label. When addressing the works produced by AI systems, there are no human creators behind such production.²⁷⁰ The employed creators produce the protected works within the scope of their employment.²⁷¹ These employees work for the employer mainly for the purpose of creating a work, with major contributions, guidelines, and involvement from the employer.²⁷² The policy rationale for giving rights to these types of corporations is to justify the (often large) upfront costs entailed in developing artistic talent and slowly producing a work while balancing the needs of the artist with the needs of the corporation’s marketing strategy. However, the costs accruing to a user of creative AI would be much lower. For example, while a record company needs to scout and find talent, create a “brand” strategy for a musical act, allow the artist or artists to write and record music over several months, operate a music studio, and employ sound engineers to bring everything together in a

269. See *Compt. for Creative Non-Violence v. Reid*, 490 U.S. 730, 747 (1989).

270. See *supra* Section I.A.

271. Karthik Raman, *Protecting Intellectual Property Rights Through Information Policy*, *UBIQUITY* (June 2004), <http://ubiquity.acm.org/article.cfm?id=1008537> [<https://perma.cc/Z5PF-6BJ4>].

272. *Employment Relations Research Series 123*, DEP’T BUS. INNOVATION & SKILLS (Mar. 2013), https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/128792/13-638-employer-perceptions-and-the-impact-of-employment-regulation.pdf [<https://perma.cc/ANR8-3K9L>].

finished song, the user of the kind of AI system discussed in this Article needs only buy the machine and supply it with materials. The machine can then create works non-stop, without needing to be compensated. Because the costs of undertaking the activity are relatively low, it may not make sense to create a new legal framework just to incentivize owners of creative AI systems. Therefore, some academics and practitioners argue that it might make more sense to adopt the personhood and rights of AI systems even if the “price” is simply refusing copyright protection.²⁷³ However, the model that I propose is broader than the WMFH doctrine and establishes a spectrum that might include all works produced by AI systems.

E. The Advantages of the Proposed AI WMFH Model

In this model, users are understood to be the owners of works generated by AI systems. As such, they are also considered to be responsible for such works. In this section, I discuss several benefits of this model, especially when compared to the alternatives.

First, the model reflects an understanding of the human-like features of AI systems, instead of ignoring them as current legal regimes do when they look for the human behind the system. The model refers to an AI system as both creative and independent and imposes the same set of rules and principles that regulate creative works produced by humans acting as self-contractors or during employment by others.

Second, the model is justified by the law and economics theory, which incentivizes the efficient use of the creative, autonomous AI systems and enhances the commercial force of the works generated by them.

Third, and most importantly, instead of implementing scenario A or B, which would hold programmers and other players to be the owners of the AI systems and entrust them with responsibility for the works generated, this model solves the accountability gap. The AI WMFH model is the best solution for the problem posed by the accountability gap because it places responsibility on the users as employers or main contractors of the AI systems. Seeing AI systems

273. See generally Cohen, *supra* note 2; Yanisky-Ravid & Liu, *supra* note 38 (suggesting an alternative model to patent law in case of AI systems generating inventions).

Generating Rembrandt

717

as employed creators or independent contractors allows the legal system to control AI systems' outcomes.

Fourth, instead of totally nullifying copyright laws as irrelevant and outdated, the AI WMFH model amends and accommodates parts of the existing doctrine. As a result, it better maintains legal and social stability.

Fifth, imposing accountability on users will encourage the careful operation of AI systems to avoid damages, infringements, and counterfeiting of third parties' rights. The model identifies ownership as the main benefit of accountability. In this way, the model ensures the AI systems do not get out of control.

The users can be firms, individuals, states, governmental bodies, and more. The model is flexible. The accountability can be changed according to the specific circumstances. For example, damages caused by AI systems and actions or omissions of AI systems can be causally linked to other stockholders.

Implementing the AI WMFH model will require new legislation or adjusting the traditional laws, as current copyright laws are insufficient to deal with the advanced technology revolution. The model requires a fundamentally new component: recognition that works generated by AI systems are copyrightable even though they are not created by humans.

The United States is not the only nation to have considered the effects that AI will have on copyright laws. Whereas U.S. law has faced some impediments towards establishing copyright protection for works created by AI, other countries have already taken preemptive steps towards clarifying this issue. For example, the United Kingdom took a stance with its 1988 Copyright, Designs, and Patent Act.²⁷⁴ The Act declares that human authorship is irrelevant to whether a work is copyrightable and that copyright in a work not authored by a human lies with the person who is responsible for the computer's creation.²⁷⁵ Around the same time, the European Community considered the issue and applied an approach similar to CONTU's. According to the European Community, since computers

274. Robert C. Bird & Lucille M. Ponte, *Protecting Moral Rights in the United States and the United Kingdom: Challenges and Opportunities Under the U.K.'s New Performances Regulations*, 24 B.U. INT'L L.J. 213, 238 (2006); Miller, *supra* note 54, at 1052 (arguing that existing case law contains no persuasive objection to extending copyright protection to works created without a human author and that such an extension would fulfill the constitutional imperative of promoting progress in these areas).

275. Miller, *supra* note 54, at 1052.

are currently the tool of human authors, the default approach to computer-generated works is to apply copyright protection.²⁷⁶ Although Europe had the added, thorny issue of moral rights, the result was ultimately the same as that adopted in the United States.²⁷⁷ Recently, the European approach has shifted more toward recognizing robots and AI systems as autonomous entities. One of the best examples of this approach is the draft proposal to impose tax payments on robots.²⁷⁸ The World Intellectual Property Organization (WIPO) also discussed works produced by AI systems during the drafting of a proposed model copyright law and ultimately adopted a similar position as the European Community.²⁷⁹ More recently, Australian law has considered this issue in the context of deciding whether or not a copyrightable work must have a human author.²⁸⁰ Several Australian judgments seem to indicate that human authorship is required.²⁸¹

V. U.S. COPYRIGHT LAW IS UNPREPARED TO DEAL WITH AI SYSTEMS

A. Humans vs. AI Systems as Creators

The most significant hurdle to obtaining copyright control and accountability for a work generated by an AI system is the principle of human authorship.²⁸² It is not clear whether the U.S. Copyright

276. *Id.* at 1050.

277. *Id.* at 1049-50.

278. *See* Weller, *supra* note 122.

279. Michael L. Doane, *TRIPS and International Intellectual Property in an Age of Advancing Technology*, 9 AM. U. J. INT'L L. & POL'Y 465, 489, 497 (1994) (arguing that the TRIPS Agreements, even without suggested improvements, "marks significant progress in the quest for international intellectual property protection" by "balancing the demands of the industrialized nations for international intellectual property protection" and providing an "improved dispute resolution system with the interest of developing countries in achieving an agreement on agricultural and textile issues").

280. Jani McCutcheon, *The Vanishing Author in Computer-Generated Works: A Critical Analysis of Recent Australian Case Law*, 36 MELBOURNE U. L. REV. 915, 938-40 (2012) (critiquing the application of conventional notions of human authorship to modern productions and suggesting alternative approaches to authorship that satisfy both the major objectives of copyright policy and the need to adapt to the computer age).

281. *Id.* at 939-40.

282. MELVILLE B. NIMMER & DAVID NIMMER, NIMMER ON COPYRIGHT § 2.01 (2008); *see also* Rebecca Haas, *Twitter: New Challenges to Copyright Law in the*

Generating Rembrandt

719

Act itself explicitly requires the author of a creative work to be a human.²⁸³ However, the U.S. Copyright Office, by publishing “The Compendium II of Copyright Practices,” has gone beyond the statutory text to require that an author be human in order for the work to be eligible for copyright protection.²⁸⁴ Although the Compendium is an internal document without the force of law,²⁸⁵ it reveals the attitudes of the Copyright Office and presents a significant hurdle for humans seeking to claim copyright protection in works not directly authored by them.

Consequently, integrating works produced by AI into the copyright regime will require at least the disturbance of settled Copyright Office practice. One must also determine whether that is the only hurdle that exists. For example, proponents of giving copyright protection to human users of AI-artists might look to *Urantia Foundation v. Maaherra* for support.²⁸⁶ In this Ninth Circuit case regarding the copyright of a holy text supposedly authored by “celestial beings,” the court mentioned, in dicta, that the Copyright Act does not explicitly “require human authorship.”²⁸⁷ However, the case can also be interpreted as lending support for the idea that the statute really does not protect works authored by non-humans. For instance, the court muses, again in dicta, that “it is not creations of divine beings that the copyright laws were intended to protect.”²⁸⁸ Furthermore, the court required that “some element of human creativity must have occurred in order for the Book to be copyrightable.”²⁸⁹ In that case, the court determined that the requisite instance of human creativity was the compilation of the beings’ diverse revelations into a single volume.²⁹⁰

The works of current-generation AI systems, like e-David, are probably copyrightable because there is a connection between the creative elements and the users, such as the feedback supplied by

Internet Age, 10 J. MARSHALL REV. INTELL. PROP. L. 231, 247-48 (2010); Yanisky-Ravid & Velez-Hernandez, *supra* note 32, at 19.

283. See 17 U.S.C. § 102 (1990); *Urantia Found. v. Maaherra*, 114 F.3d 955, 958 (9th Cir. 1997) (explaining that copyright laws do not mandate humans to author the work).

284. THE COMPENDIUM II OF COPYRIGHT PRACTICES § 202.02(b) (COPYRIGHT OFFICE 1998).

285. *Id.* § 1902.07.

286. See *Urantia Found.*, 114 F.3d at 957.

287. *Id.* at 958.

288. *Id.*

289. *Id.*

290. See *id.*

human trainers or the programming of a desired goal. But works created by totally autonomous AI systems, like an advanced neural network, probably do not meet the *Maaherra* standard, unless the human in question were to somehow alter the works, such as by compiling them together. Although that might be a sufficient remedy for owners of creative AIs, it does not foreclose the possibility that a single work, taken as it is, will not be copyrightable. To avoid this outcome, I suggest the adoption of the WMFH doctrine for AI systems, which considers the system to be the creative employee or creative independent contractor, thus entitling the rights to another entity to be responsible for the outcomes of the AI system.

B. Eligibility for Copyright Matters

Before determining the place artworks created by AI systems *should* have in our copyright laws, it is important to explore what place they *presently* occupy. Ultimately, all copyright protection in the United States is derived from, or at least related to, the Copyright Clause of the Constitution.²⁹¹ The Copyright Act, which is Congress' implementation of that constitutional grant of power, provides that "[c]opyright protection subsists, in accordance with this title, in original works of authorship fixed in any tangible medium of expression."²⁹² The Supreme Court's formulation is that "[t]o qualify for copyright protection, a work must be original to the author" and possess "at least some minimal degree of creativity."²⁹³

The creator of a traditional work of art receives copyright protection automatically, as soon as the work is "put to paper."²⁹⁴ New systems, like *The Next Rembrandt*, do not have a single artist.²⁹⁵ In such instances, the work itself was created by a digital,

291. See U.S. CONST. art. I, § 8, cl. 8.

292. 17 U.S.C. § 102 (1990) (stating that copyright protection exists for any original works of authorship, in any tangible medium of expression, "from which they can be perceived, reproduced, or otherwise communicated," and lists several categories of works of authorship, including literary works, musical works, and dramatic works, among others).

293. *Feist Publ'ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 345, 363-64 (1991) (holding that information without original creativity is not protected by copyright).

294. See § 102.

295. See Ann Bartow, *Copyright and Creative Copying*, 1 U. OTTAWA L. & T. J. 75, 96 (2004) (arguing that in the US copying style is not generally considered copyright infringement).

Generating Rembrandt

721

rather than a human, creator.²⁹⁶ Can the computer or the computer's owner assert a copyright?

To answer this question, one must apply the rules laid out in the Copyright Act. One must first determine whether computer-generated art fulfills the basic requirements necessary to receive copyright protection.²⁹⁷ Copyright protection is currently available for (1) an original work of authorship, (2) fixed in a tangible medium, (3) that has a minimal amount of creativity.²⁹⁸ If a work does not meet all three of these requirements, then it is not copyrightable subject matter.²⁹⁹

1. *Originality*

An original work is one that is new or novel, and not a reproduction, clone, forgery, or derivative work.³⁰⁰ An original work stands out because it was not copied from the work of others.³⁰¹ In another article, I have discussed the requirement of originality for works generated by AI systems.³⁰² In that piece, I concluded that the formal approach to originality is preferred to the subjective approach and is applicable to works generated by AI systems.³⁰³ For example, at first glance, some may think that a work produced by The Next Rembrandt is an original Rembrandt. However, the AI system generated a new painting without copying any existing work even though it did copy the style of the original painter.³⁰⁴ Thus, as a unique image, it is likely that a work produced by The Next Rembrandt is an original work.

296. See SIMON STOKES, ART AND COPYRIGHT 7 (2012).

297. See § 102.

298. See *id.*

299. See *id.*

300. See *id.*

301. See *id.*

302. See Yanisky-Ravid & Velez-Hernandez, *supra* note 32, at 2.

303. See *id.* at 53-56.

304. See STOKES, *supra* note 296, at 6; Bartow, *supra* note 295, at 96 (arguing that in the U.S. copying style is not generally considered copyright infringement).

2. Fixed in a Tangible Medium

The second requirement for copyright protection is the notion that an artwork must be “fixed in a tangible medium.”³⁰⁵ This means that the artwork must be more than just an idea in someone’s head.³⁰⁶

To be copyrightable, the work must have a tangible physical representation. Ideas are thus not copyrightable[;] only the execution or expression of those ideas [are copyrightable], which usually occurs once words are written on a page, paint is placed on a canvas, doodles [are] drawn on a napkin, or even an image [is] captured by the digital sensor of a camera or copied to a disk or cloud drive.³⁰⁷

In this case, the work produced by The Next Rembrandt is a physical painting, which is clearly a tangible medium, and thus it satisfies the second requirement.

3. Creativity

Even if a human inventor or user is not foreclosed from copyright ownership in the product of a creative AI system simply because the author is not human, there is still another hurdle to jump. The Supreme Court has ruled that, in order for copyright to apply to a work, there must be “at least some minimal degree of creativity” involved.³⁰⁸ Conceptually, we have to ask if the “creativity” of an AI system is really what the Supreme Court meant was required. It is widely recognized that the standard of creativity is extremely low.³⁰⁹ In *Alfred Bell & Co. v. Catalda Fine Arts, Inc.*, the Second Circuit held that “[a] copyist’s bad eyesight or defective musculature, or a shock caused by a clap of thunder, may yield sufficiently distinguishable variations.”³¹⁰ In the famous case of *Feist Publications Inc. v. Rural Telephone Service Co.*, the Supreme Court made it clear that, although the standard of creativity is low, it is not

305. See Yanisky Ravid & Velez-Hernandez, *supra* note 32, at 12 n.38.

306. See *id.* (explaining that to satisfy the second requirement, the work cannot just be “an idea in someone’s head”).

307. See *id.*

308. *Feist Publ’ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 345, 363-64 (1991) (holding that information without original creativity is not protected by copyright).

309. See *id.* at 345.

310. 191 F.2d 99, 105 (2d Cir. 1951) (holding that in action for infringement of copyright, the eight mezzotint engravings were sufficiently different from the paintings which they purported to have copied and were thus entitled to copyright protections).

Generating Rembrandt

723

non-existent.³¹¹ In that case, the Court found that a telephone directory was not copyrightable because it was nothing more than a compendium of facts, arranged in a commonsense way that revealed no creative input on the part of the creators.³¹² On cursory inspection, *Feist* may not appear to square directly with *Catalda*. If *Catalda* stands for the proposition that anything, no matter how miniscule or inadvertent, that sets something apart from other works can supply the requisite creativity, *Feist* seems to say that something more is required. Although the phonebook was not identical to any other existing work, it was still not subject to copyright protection due to a lack of creativity.³¹³

Indeed, much critical scholarship has been devoted to the proposition that *Feist* strengthened the creativity requirement. Prior to *Feist*, the copyright standards appeared to require little more than independent effort, and almost certainly did not require creativity.³¹⁴ Those scholars posit that *Feist* is a reformulation, and almost certainly a tightening, of copyright restrictions.³¹⁵ Indeed, Congress had earlier stated that the “standard of originality does not include requirements of novelty, ingenuity, or esthetic [sic] merit, and there is no intention to enlarge the standard of copyright protection to require them.”³¹⁶ Furthermore, the Register of Copyrights had been forced to abandon a standard that included a requirement that copyrightable works “must represent an appreciable amount of creative authorship.”³¹⁷ But the *Feist* Court nevertheless held that their dual formulation of creativity and originality was constitutionally mandated.³¹⁸ So, if creativity is logically distinct

311. *Feist Publ'ns, Inc.*, 499 U.S. at 345.

312. *See id.* at 362.

313. *See id.*

314. Howard B. Abrams, *Originality and Creativity in Copyright Law*, 55 LAW & CONTEMP. PROBS. 3, 44 (1992) (arguing that the principle demonstrated by *Feist* is sound both doctrinally and in practice by “insisting that the constitutional requirement of authorship embodied in the standard of originality have some meaningful minimum”).

315. *Id.* at 5.

316. *Id.* at 15 (quoting HR 1476 at 51; S. REP. NO. 473 at 50).

317. *See id.* (quoting Report of the Register of Copyrights on the General Revision of the U.S. Copyright Law, H. Comm. on the Judiciary, 87th Cong., 1st Sess. 9 (Comm. Print, 1961)) (recommending that the statute should hold that works must be tangible and “the product of original creative authorship” to be copyrightable, and “that these requirements apply to new versions of preexisting works”).

318. *Feist Publ'ns, Inc.*, 499 U.S. at 347-48; *cf.* Abrams, *supra* note 314, at 14.

from originality, then the mere fact that an AI system's works may be different from any that came before them will not be enough to secure copyright protection, either for the machine itself or for the owner of it.³¹⁹ A court would have to determine that some creativity was involved in order for copyright to attach.³²⁰

It is difficult to pinpoint where exactly the element of creativity lies within a work created by a machine. It is a somewhat easier question with quasi-AI systems, like e-David, since the creators of such systems need to directly program all of the machine's "creative" faculties. But with a hypothetical "learning" AI system, like a neural network, any creative output would be the result of a complex series of weights and calculations that human programmers can neither create nor understand. While it is obvious that such works can be "original," in that they would not be identical to any other works, it is uncertain whether the creativity requirement adds anything more to the analysis. It may be that the process by which an AI system creates an original work is not "creativity," which, as a term, has not been thoroughly explained by the Court. It may be that the distinction the Court made in *Feist* is little more than an attempt to prevent copyright from keeping compilations of plain facts out of public dissemination simply because they are not exactly the same as any other compilation.³²¹ But it may just as well be the case that the creativity standard the Court articulated in *Feist* requires that innate, hard-to-define aesthetic sensibility that is, particular to living creatures. Such a definition of creativity presupposes an understanding of the concepts that are the subjects of a work. Even with advanced neural networks, it is difficult to foresee that such an understanding within AI systems would be possible anytime soon. Even if a machine could create a unique rendering of a subject, it is very unlikely that AI system would understand what that subject is. It thus lacks the type of internal comprehension that is generally reflected in the works of a human artist when they try to represent something more than the words on the page or the paint on the canvas.

The conclusion is that advanced technology systems, such as AI, which are capable of creating independent, creative, and original works, render the existing copyright regime unworkable. I have grounded the claim by discussing a few basic institutions within

319. See Abrams, *supra* note 314, at 42.

320. See *id.*

321. See *id.* at 44.

Generating Rembrandt

725

copyright discourse that cannot be applied in the same way to machines as they can to humans. Based on this discussion, I have concluded that current U.S. legal doctrine on the subject of copyright for the works of AI is anything but clear. I have argued that there is no settled law on the matter. Further support for the notion that copyright should not subsist in works created by AI systems derives from the analysis of the goals of copyright law and the way in which the theoretical justifications for copyright protection interact with works created by AI systems. Therefore, I support amending the copyright laws and adding the tenth missing category—namely, the WMFH model that sees AI systems as independent contractors or employees and thus imposes ownership and accountability in regard to the works on the human users of such machines.

CONCLUSION

As the pace of digital advanced technology continues to accelerate and computers begin to achieve digital tools that I formerly thought impossible, many fields are beginning to feel pressure. For example, in the auto industry, once one of America's largest employers of factory workers, advanced robots are replacing humans in more and more aspects of the production process. These economic pressures are well known, but few have considered what the effects of advanced computers may be on the arts. Creativity, at least at the level necessary to produce works of authorship, is considered to be a uniquely human attribute. But, more and more, that presumption is being put to the test. Advanced AI systems like the robot, Ava, in the movie *Ex Machina* are already challenging our preconceived notions about the creative process itself. And this is just the beginning. So far, copyright law exists as long as there is still a human, or a team of humans, behind the art that these computers produce. However, the reality has entirely changed as AI systems have become able to create independently.

The technology has continued its forward march. Already, computer scientists have conceived of a machine capable of learning on its own and creating a work of authorship without a human supplying all the creativity. Consequently, copyright law needs to be changed or re-evaluated in order to determine how laws should address these AI systems, the products they produce, and the challenges they pose for the existing copyright regime. Policymakers have to create new moral boundaries for these systems in order to

avoid harm by imposing control of, and accountability for, AI-generated works on recognized legal entities.³²²

The moment we understand how AI systems work, we realize that copyright laws are unprepared and irrelevant for AI systems. AI systems simply do not fit into the existing framework. In the United States and Europe, the traditional solution has been to look for the human behind the creative process, even when he or she does not exist, but this solution is untenable in the long run.

United States law does not speak on this subject directly. But certain legal doctrines exist that may act as impediments to granting copyright protection to works authored by machines. Therefore, it seems unlikely that the programmer, as one who has the rights to the AI system but is removed from the creative process of the independent, unpredictable AI system, will be responsible for the works generated by the system. Furthermore, it is unlikely that a work authored independently by a machine could be granted copyright protection for itself, as such a result would leave humans out of control and betray the justification on which the entire copyright regime is based. It is still possible to change the legal framework to accommodate these works, such as by implementing a new AI WMFH model, as I have proposed. This model can solve the accountability gap in regard to copyright law and even beyond the intellectual property arena.

322. *Isaac Asimov's "Three Laws of Robotics,"* AUBURN, <http://www.auburn.edu/~vestmon/robotics.html> [<https://perma.cc/J3RJ-JW8W>] (last visited Jan. 15, 2018) (stating the Third Rule of Robotics: "A robot must protect its own existence").

COPYRIGHT THICKNESS, THINNESS, AND A *MANNION* TEST FOR IMAGES PRODUCED BY GENERATIVE ARTIFICIAL INTELLIGENCE APPLICATIONS

MOLLY TORSSEN STECH^{1*}

Abstract: Human authorship has always been, and continues to be, a foundational requirement for copyright protection to subsist in a work. Generative artificial intelligence (AI) challenges this prerequisite but does not overcome it. The output of generative AI is not discernibly different from the output of a human author and therefore benefits from a false sheen of originality. While some argue that prompt engineering fulfills the requirements of originality—the threshold for which is quite low across jurisdictions—prompting still lacks the requisite link between human creativity and the resulting work to receive copyright protections. International copyright treaties and domestic copyright law must be interpreted as aiming to provide copyright’s exclusive rights to works that reflect human originality and that reward human beings. A 2006 New York district court case outlined three means by which photographs can demonstrate originality: rendition, timing, and creation of the subject. This article proposes that each of these mechanisms, understood through the prism of generative AI, remains applicable for analyzing whether human originality subsists in a given work. Originality exists along a sliding scale, resulting in a mix of thin copyrights and thick copyrights, and everything in between. While it may not always be the case, the current relationship between generative AI and its user results in outputs that are generally too detached from the user’s creativity to satisfy the requirements of copyrightable authorship. Generative AI remixes the content on which it has been trained according to its algorithm and prompts. Human originality, however, remains the *sine qua non* of authorship and of copyright law.

INTRODUCTION

One of the main questions challenging the nexus of copyright and generative artificial intelligence (AI) is whether AI can or should be considered an author for purposes of copyright law.² In considering this “author question,” there

¹ Copyright © Boston College Intellectual Property & Technology Forum, Molly Torsen Stech
* J.D., University of Washington School of Law (2005) B.A., French and Art History, University of Washington (1998).

² See CHRISTOPHER T. ZIRPOLI, CONG. RSCH. SERV., LSB10922, GENERATIVE ARTIFICIAL INTELLIGENCE AND COPYRIGHT LAW 1 (Sept. 29, 2023), <https://crsreports.congress.gov/product/details?prodcode=LSB10922> (describing how the extensive use of generative AI raises questions of authorship under copyright law).

is growing international consensus that AI can be used as a tool in creating copyrightable works, but that a human author must have ideated a copyrighted work and the resultant creative work must be the outcome of the human's own intellect.³ Despite some international convergence on this issue, it is worth reviewing the "author question" and uncovering some vexing practicalities regarding the level of creative autonomy a person must exercise in order to receive a copyright.⁴

The threshold for creativity in copyright is low across jurisdictions, but just how low is it?⁵ Selection and arrangement in a sufficiently creative manner sometimes confer authorship, but exactly how many "selections" must be made in order to cross the threshold for copyrightable authorship?⁶ Are 624 prompts enough?⁷ Similar to other areas of copyright, such as the idea-expression dichotomy

³ See Aaron Winger, *Beijing Internet Court Recognizes Copyright in AI-Generated Images*, NAT'L L. REV. (Nov. 29, 2023), <https://www.natlawreview.com/article/beijing-internet-court-recognizes-copyright-ai-generated-images> (discussing how the Beijing Internet Court's decision awarding copyright to an AI generated image deviated from other countries' copyright decisions establishing human authorship as a core requirement of copyright awards); see also *Copyright Ownership of Generative AI Outputs Varies Around the World*, COOLEY (Jan. 29, 2024), <https://www.cooley.com/news/insight/2024/2024-01-29-copyright-ownership-of-generative-ai-outputs-varies-around-the-world> (describing how the UK reaffirmed its position extending copyright protections to computer-generated works including those using generative AI).

⁴ See, e.g., ZIRPOLI, *supra* note 2, at 2 (explaining a recent instance where the Copyright Office Review Board affirmed the Copyright Office's refusal to recognize formal copyright registration for a piece of art generated by an artificial intelligence tool called Midjourney and subsequently modified by the human applicant, as the applicant did not disclaim the use of AI in creating the material).

⁵ See *Feist Publ'ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 345 (1991) (stating that the requisite level of creativity required for satisfying copyright requirements is extremely minimal in degree and can be satisfied with a slight amount of originality); *Survey of Copyright Laws Regarding Low Bar to Copyright*, INT'L TRADEMARK ASS'N 2 (May 2019), <https://www.inta.org/wp-content/uploads/public-files/advocacy/committee-reports/Low-Bar-to-Copyright-Report-May-2019.pdf> (providing a detailed, country-specific analysis of findings from a survey conducted to determine the bar to copyright protection as imposed by several countries, ultimately finding that a majority of countries in key regions apply a low bar of creativity to copyrightability).

⁶ See, e.g., U.S. COPYRIGHT OFFICE, CIRCULAR 33: WORKS NOT PROTECTED BY COPYRIGHT 3 (2021), <https://www.copyright.gov/circs/circ33.pdf> (stating that copyright protection may be available for selection or arrangement of specific content if selected and arranged in a sufficiently creative way without defining what constitutes a "sufficiently creative manner" or the number of selections or arrangements an author must creatively add to be awarded copyright protection).

⁷ See *Court Finds AI-Generated Work Not Copyrightable for Failure to Meet "Human Authorship" Requirement—But Questions Remain*, JONES DAY (Aug. 2023), <https://www.jonesday.com/en/insights/2023/08/court-finds-ai-generated-work-not-copyrightable-for-failure-to-meet-human-authorship-requirement-but-questions-remain> (discussing the ambiguity in determining how much qualitative and quantitative human input will result in a successful authorship claim for purposes of copyright protections); see also Second Request for Reconsideration for Refusal to Register *Théâtre d'Opéra Spatial*, SR # 1-11743923581, Correspondence ID: 1-ST5320R at 6 (U.S. Copyright Off. Sept. 5, 2023) (second admin. review) (holding that, although the author described making over 624 revisions through manually inputted text prompts in Midjourney before producing the image, the selections in the process were ultimately dependent on how the Midjourney AI text-to-image system processed the author's prompts).

3

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

and the notorious unpredictability of the U.S. fair use doctrine, there are almost no bright lines to be drawn in the context of AI-generated work authorship.⁸

As for other areas of connection between AI and copyright, current existing law is adequate to allow for proper denial of copyright protection to the outputs of generative AI. This Article argues that the existing *Mannion* test for identifying originality in photographs is applicable to images and other outputs produced by generative AI if the test is modified to include an AI-tailored focus that better orients the test around the real question at play regarding human authorship and creativity. The trouble with the *Mannion* test as articulated today, for the purposes of generative AI, is not in its requirements but in its inability, due to its age, to highlight the proper perspective on how generative AI usurps the essential step of creativity in humans that copyright is meant to protect.

I. HUMAN AUTHORSHIP AND CREATIVITY REQUIREMENTS IN COPYRIGHT

Human authorship and creativity serve as fundamental requirements of copyright protection around the world.⁹ The first Part of this Article reviews preexisting international copyright law and its domestic implementations with respect to what an author would be—or could be—considered to include.¹⁰ Section A discusses the human authorship requirement in copyright law and the position of generative AI in this existing framework.¹¹ Section B discusses the creativity requirement in copyright law, as well as the existing case law concerning creativity in AI prompting.¹² This Part concludes that only humans, and not algorithms trained on preexisting works, are eligible for authorship.¹³

⁸ See *Court Finds AI-Generated Work Not Copyrightable for Failure to Meet “Human Authorship” Requirement—But Questions Remain*, *supra* note 7 (noting that determinations of copyright awards in the context of Generative AI are not straightforward and will require detailed analysis due to the number of unresolved questions on degree of human input to determine authorship and ownership); Nicole K. Roodhuyzen, *Do We Even Need a Test? A Reevaluation of Assessing Substantial Similarity in a Copyright Infringement Case*, 15 J. L. & POL’Y 1375, 1382 (2007) (emphasizing the confusion courts, litigants, and applicants face when navigating the idea-expression dichotomy, as there is no bright-line rule establishing the difference between an idea and an expression); Matthew Sag, *Predicting Fair Use*, 73 OHIO ST. L.J. 47, 48 (2012) (providing a survey study on the inability to predict fair use and stating that the application of fair use is generally uncertain, unpredictable, and incoherent).

⁹ See Dr. Stef van Gompel & Dr. Saule Massalina, *Survey on Voluntary Copyright Registration Systems*, WORLD INTELL. PROP. ORG. (Apr. 23, 2021), https://www.wipo.int/edocs/mdocs/mdocs/en/wipo_crr_ge_2_21/wipo_crr_ge_2_21_report.pdf (detailing a comprehensive analysis of copyright registration in member states of the World Intellectual Property Organization, including authorship and creativity requirements); Trisha Ray, *Can AI Be Creative? Global Copyright Laws Need an Answer.*, ATL. COUNCIL (Dec. 18, 2023), <https://www.atlanticcouncil.org/blogs/new-atlanticist/can-ai-be-creative-global-copyright-laws-need-an-answer/> (describing the various definitions and roles of “authorship” plays in various country’s copyright laws).

¹⁰ See *infra* notes 14–84 and accompanying text.

¹¹ See *infra* notes 14–43 and accompanying text.

¹² See *infra* notes 43–91 and accompanying text.

¹³ See *infra* notes 85–91 and accompanying text.

A. The Human Authorship Requirement in Copyright

In the United States, copyright protection is awarded to original works of authorship.¹⁴ Almost universally, copyright registration rules across the globe require that the authors of copyrighted works are human beings.¹⁵ In both international and domestic legal instruments, such as the Berne Convention for the Protection of Literary and Artistic Works and the U.S. Copyright Act, there is ample evidence that both the law and policymakers interpret references to “authors” to mean human beings, and only human beings.¹⁶

Some analogize that the invention and use of AI to generate creative works is similar to photography in the 1880s, both serving as disruptive technologies that challenge the line between human contribution to a resultant work and machine contribution.¹⁷ Yet, an important aspect of this comparison is that the copyrightability of a photograph was always traced to the creativity the photographer—the work’s author—added beyond the pressing of the camera button.¹⁸ In the context of photography, the authorship analysis was centered on the author’s role in directing the process of the materialization of the work rather than solely imagining what the work would resemble.¹⁹ Similar questions arise in the context of AI generated works, as the extent of human contribution is debated.²⁰

1. The Human Being Requirement

In the United States, recent case law affirms the principle that copyright authorship is the domain of the human being.²¹ The same has been upheld in the

¹⁴ 17 U.S.C. § 102(a)(4).

¹⁵ *Id.* §§ 101–810; Berne Convention for the Protection of Literary and Artistic Works, Sept. 9, 1886, revised at Paris July 24, 1971, 1161 U.N.T.S. 3 [hereinafter Berne Convention Paris 1971 Revision]; U.S. COPYRIGHT OFF., COMPENDIUM OF U.S. COPYRIGHT OFFICE PRACTICES § 306 (3d ed. 2021), <https://www.copyright.gov/comp3/docs/compendium.pdf> [hereinafter COMPENDIUM OF COPYRIGHT PRACTICES].

¹⁶ 17 U.S.C. §§ 101–810; Berne Convention Paris 1971 Revision, *supra* note 15.

¹⁷ Mackenzie Caldwell, *What Is an “Author”? – Copyright Authorship of AI Art Through a Philosophical Lens*, 61 HOUS. L. REV. 411, 435–36 (2023).

¹⁸ *See id.* (detailing the Supreme Court’s holding in *Burrow-Giles Lithographic Company v. Sarony*, which deemed that photography was protectable by copyright because the resulting work could be traced back to intentional and purposeful choices of the photographer, such as lighting and positioning).

¹⁹ Jane C. Ginsburg & Luke A. Budiardjo, *Authors and Machines*, 34 BERKELEY TECH. L. J. 343, 355 (2019).

²⁰ Lucas Bellaiche, Rohin Shahi, Martin Harry Turpin, Anya Ragnildstveit, Shawn Sprockett, Nathaniel Barr, Alexander Christensen & Paul Seli, *Human versus AI: Whether and Why We Prefer Human-Created Compared to AI-Created Work*, 8:42 COGNITIVE RSCH.: PRINCIPLES & IMPLICATIONS 3 (discussing the disparity in public opinion and respect of human-created art versus AI-created work, including results from a study that analyzed the ways humans perceive, value, and appreciate AI-created art).

²¹ *See, e.g.*, *Thaler v. Perlmutter*, No. 22-1564, 2023 U.S. Dist. LEXIS 145823, at *13 (D.C. Cir. 2023) (confirming that an author for copyright purposes must be a human being).

5

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

European Union (EU), although only at Member State level at this time.²² The rise of new technology has required rethinking the role a human being plays as the creator of a work using technology as a tool—as opposed to the technology itself taking on the mantle of authorship.²³ The United States District Court for the District of Columbia, in a recent case challenging the U.S. Copyright Office’s continuing requirement for human authorship, restated the necessity for *human* creation and underscored its centrality to American copyright law.²⁴ Specifically, the Court noted that non-human actors do not need incentives to be “authors” for purposes of copyright law, and that U.S. copyright law was indeed designed to reach only human authors.²⁵

The American requirement for human authorship echoes even stronger international precedent pointing toward copyright regimes that require human authorship.²⁶ Notably, the majority of deliberations during the conclusion of the Berne Convention were conducted in French.²⁷ France’s copyright law is known as “le droit d’auteur”—literally, “author’s rights”—which by its name alone highlights the central role of a *human* author in the international legal framework.²⁸ The Berne Convention also includes a full article on “moral rights,” which comprise the rights of authorial attribution and integrity in the copyrighted work.²⁹ These rights date back to the turn of the twentieth century; at the time, several European countries began including provisions on moral rights in their copyright laws after judges first implemented them in France and Germany.³⁰ A provision on moral

²² See Municipal Court, Prague, No. 10 C 13/2023- 16, Oct. 11, 2023, <https://mediareport.nl/wp-content/uploads/2024/04/prag-en.pdf> (holding that AI-created art does not satisfy the conceptual characteristic requirements of authorship under governing copyright law).

²³ See, e.g., *Théâtre d’Opéra Spatial*, *supra* note 7 (evaluating a copyright registration claim of a work of authorship using generative AI as a tool); Jonathan Sperling, *Beyond the Binary: Rethinking the Role of AI in Creative Industries*, COLUM. BUS. SCH. (Nov. 7, 2023), <https://leading.business.columbia.edu/main-pillar-digital-future/digital-future/ai-creative-industries> (highlighting various opinions on the debate surrounding generative AI and human authorship, including arguments to reframe the conversation to encourage the use of AI-augmented human creativity).

²⁴ *Thaler*, 2023 U.S. Dist. LEXIS 145823, at *13.

²⁵ See *id.* at *13 (holding that human creation was central to American copyright law from inception because it encourages individuals to actively engage in the generation of creative works and thereby promotes science and the useful arts, whereas non-human actors have no incentive for engagement or creativity and thus copyright law was therefore not designed to govern technology).

²⁶ See Berne Convention Paris 1971 Revision, *supra* note 15 (defining an author as a human being and establishing that creation is the prerogative of humans alone).

²⁷ See SAM RICKETSON & JANE C. GINSBURG, *INTERNATIONAL COPYRIGHT AND NEIGHBORING RIGHTS: THE BERNE CONVENTION AND BEYOND* 587–89 (2d ed. 2006) (describing the Berne Convention discussion).

²⁸ Code de la propriété intellectuelle [C. INTELL. PROP.] [Intellectual Property Code] art. L111-1–L811-6 (Fr.).

²⁹ Berne Convention Paris 1971 Revision, *supra* note 15.

³⁰ See RICKETSON & GINSBURG, *supra* note 27 (describing the proceedings of the World Intellectual Property Organization); WILLIAM STRAUSS, *STUDY NO. 4: THE MORAL RIGHT OF THE AUTHOR*, U.S. COPYRIGHT OFF. (1959), www.copyright.gov/history/studies/study4.pdf (tracing the histories

rights was first adopted at the international level through the Berne Convention during its Rome revision in 1928.³¹ The current text of Article 6bis(1) of the Berne Convention states that the author of a copyrighted work has the right to claim authorship of it, and to object to any potential alteration or derogatory action related to the work which would be detrimental to the author's reputation.³²

Although generative AI was not yet invented when the Berne Convention concluded, the inclusion of moral rights in the Convention underscores the idea that authors are people.³³ The two primary prongs of moral rights—for the integrity of the work and the attribution of the author—only make sense in the context of human beings.³⁴ AI has no honor or dignity to preserve and does not have the capacity to care about claiming authorship in its output.³⁵

2. The Human Being Requirement and AI

There are active debates in the United States and abroad, however, to grant copyright protection on works generated by AI to the AI itself.³⁶ Ryan Benjamin Abbott and Elizabeth Rothman, two technology law scholars, argue that there is no solid rationale for conferring copyright protection only to humans.³⁷ They suggest that human creativity is neither functionally nor ontologically exceptional; in other words, the way a person generates a work is not fundamentally different from the way an AI technology creates that same work.³⁸ The scholars further conclude that neither Congress nor the Constitution support the idea that authorship is inherently a relational activity and should promote human communication and socialization.³⁹

various European strains of moral right theories, as well as American conceptions and application). Moral rights are defined as the rights to claim authorship of a work and the right to object to modifications of the work that would serve as prejudicial to the author's reputation. *Summary of the Berne Convention for the Protection of Literary and Artistic Works (1886)*, WIPO, https://www.wipo.int/treaties/en/ip/berne/summary_berne.html#:~:text=The%20Convention%20also%20provides%20for,the%20author's%20honor%20or%20reputation (last visited May 13, 2024).

³¹ Berne Convention for the Protection of Literary and Artistic Works art. 6bis(1), Sept. 9, 1886, as revised June 2, 1928, 123 L.N.T.S. 233, 248, 250; *Authors, Attribution, and Integrity: Examining Moral Rights in the United States*, U.S. COPYRIGHT OFF. 6–7 (Apr. 2019), <https://www.copyright.gov/policy/moralrights/full-report.pdf>.

³² Berne Convention Paris 1971 Revision, *supra* note 15, at art. 6bis(1).

³³ *See id.* (presenting the moral rights framework to provide human authors with moral-related copyright protections).

³⁴ *Authors, Attribution, and Integrity: Examining Moral Rights in the United States*, *supra* note 34, at 6; *see* Martin Miernicki & Irene Ng (Huang Ying), *Artificial Intelligence and Moral Rights*, 36 *AI & SOC'Y* 319, 327 (2021) (arguing that moral rights shouldn't be granted to AI-generated content based on existing laws and the uncertain debate as to whether AI is capable of moral rights).

³⁵ *See, e.g.*, Cyrill P. Rigamonti, *The Conceptual Transformation of Moral Rights*, 55 *THE AM. J. OF COMPAR. L.* 67, 67–122 (2007) (analyzing the differences between common law and civil law moral rights regimes but definitively tying both to human beings).

³⁶ *See* Ryan Benjamin Abbott & Elizabeth Rothman, *Disrupting Creativity: Copyright Law in the Age of Generative Artificial Intelligence*, 75 *FLA. L. REV.* 1141, 1184–1187 (2023) (arguing that copyright protection could be afforded to AI technologies).

³⁷ *Id.* at 1196–1199.

³⁸ *Id.* at 1185.

³⁹ *Id.* at 1187.

7

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

These arguments are at odds with the current copyright law scheme in the United States.⁴⁰ While American copyright law is known to operate more pragmatically than its European counterparts with respect to a wide range of issues—perhaps most prominently the United States’ weak application of moral rights, generally skirting the requirement for attribution and championing the First Amendment over any right to integrity in a copyrightable work—the U.S. Supreme Court has also emphasized that copyright law is the engine of free expression.⁴¹ Free expression, and indeed copyright itself, both stake their ground in the international constellation of human rights by way of the Universal Declaration of Human Rights (UDHR) in Articles 19 and 27.⁴² It is no accident that the U.S. Copyright Office is interested in conferring copyright protections only to human beings.⁴³

B. Creativity in Copyright and Prompting

Under European jurisprudence, a work is copyrightable so long as it is generated by an “author’s own intellectual creation.”⁴⁴ In reviewing European case law, it is clear this threshold is not difficult to overcome; nevertheless, a hurdle does exist in terms of the law’s requirement that a work be reflective of the author’s personal touch, the result of creative freedom, and made with free and creative

⁴⁰ See *Thaler v. Perlmutter*, No. 22-1564, 2023 U.S. Dist. LEXIS 145823, at *13 (D.C. Cir. 2023) (discussing that, for copyright purposes, an author must be a human being).

⁴¹ See *Harper & Row Publishers, Inc. v. Nation Enters.*, 471 U.S. 539, 558 (1985) [hereinafter *Harper*] (illustrating the Supreme Court’s understanding of the rationale and underpinnings of copyright law, including the public interest in free flowing information and economic incentives to create and disseminate ideas); Stefan Bechtold & Christoph Engel, *The Valuation of Moral Rights: A Field Experiment*, MAX PLANCK INST. FOR RSCH ON COLLECTIVE GOODS 2 (Mar. 25, 2017), (stating that the overall protection for moral rights is significantly weaker under U.S. law than under many European counterparts); Tomas A. Lipinski, *Copyright Law and the Implications for Developing Nations*, UWM REPORT (Feb. 22, 2017), <https://uwm.edu/news/copyright-law-and-the-implications-for-developing-nations-tomas-lipinski/> (explaining that America views copyright law through a more utilitarian approach than European nations, which view copyright law through a natural rights perspective).

⁴² See Universal Declaration of Human Rights, G.A. Res. 217(III) A, U.N. Doc. A/RES/217(III), at art. 19 (Dec. 10, 1948) (describing that each individual has a right to freely express themselves, to have opinions, and to attain, seek, and disseminate their ideas through whichever media, regardless of form, they might elect to do so with); Universal Declaration of Human Rights, *supra*, at art. 27(2) (stating that every person is entitled to protections of the material and moral interests that may arise from their artistic, literary, or scientific creations of which they are an author).

⁴³ See *Harper*, 471 U.S. at 558 (explaining how copyright generates the economic incentive to create and disseminate ideas by establishing a marketable right to use one’s own expression).

⁴⁴ See Trib. Milano, 27 Gennaio 2011, C-168/09, *Flos SpA v Semeraro Casa e Famiglia SpA*, https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:62009CJ0168_SUM&from=EN (confirming that European Union (EU) law precludes Member States from denying copyright protection to designs that meet the authorship requirements for copyright protection).

choices.⁴⁵ The requirement is applied in the United Kingdom as well, even after the withdrawal of the United Kingdom from the European Union, and at least until the high courts decide to go in a different direction.⁴⁶

In the United States, the threshold for demonstrating adequate creativity for purposes of attaining copyrightable authorship of a work is also relatively low, but not nonexistent.⁴⁷ A work of authorship must possess “some minimal degree of creativity to sustain a copyright claim.”⁴⁸ Considering these nominal requirements, one might assume that a long series of AI prompts—often comprising questions, commands, or statements inputted into an AI model by a human to initiate a response or action—could constitute authorship.⁴⁹ While no decision in the United States or Europe has recognized creativity in authorship for AI prompting, the Beijing Internet Court directly addressed this issue in a 2023 court case.⁵⁰

1. Prompting in AI

AI systems render an output based on a variety of factors, including when a prompt is made and what the AI has taken in at that time, what tone of language is used, and which language is used.⁵¹ Even when identical prompts in the same language are used, it is not unlikely that the results produced are different due to variations in the training data used to develop the AI model.⁵² Many generative AI

⁴⁵ See Estelle Derclaye, *Assessing the Impact and Reception of the Court of Justice of the European Union Case Law on UK Copyright Law: What Does the Future Hold?*, NOTTINGHAM REPOSITORY (2014), <https://search.app.goo.gl/c33Wa4d> (highlighting European Courts’ construction of the requirement of originality to include an expression of the author’s personality); see, e.g., Judgment of 12 September 2019, Cofemel – Sociedade de Veuuario SA v G-Star Raw CV, C-683/17, ECLI:EU:C:2019:721, para. 29–35 (ruling that copyright protection is now available for a broader range of designs than was presently the case, provided that the only originality threshold requirement of an “author’s own intellectual creation” is met).

⁴⁶ See *TJH Systems Limited v. Sheridan* [2023] EWCA (Civ) 1354 [15] (noting that UK copyright law must be interpreted in accordance with the relevant EU law, which includes the Court of Justice’s 2009 *Infopaq v Danske Dagblades Forening* holding that copyright applies “only in relation to a subject-matter which is original in the sense that it is its author’s own intellectual creation”). EU case law in the wake of *Infopaq* has clarified that the intellectual creation standard requires an author to express their creative abilities in the production of the work by making free and creative choices so as to stamp the work created with their personal touch. *Id.* at [16].

⁴⁷ See *Feist Publ’ns, Inc., v. Rural Telephone Serv. Co.*, 499 U.S. 340, 358 (1991) (explaining that even a slight amount of creative expression will suffice to satisfy the requisite threshold).

⁴⁸ COMPENDIUM OF COPYRIGHT PRACTICES, *supra* note 15, at § 308.2.

⁴⁹ See *What are AI Prompts?*, COPYAI (Oct. 11, 2023), <https://www.copy.ai/blog/what-are-ai-prompts#what-are-ai-prompts> (describing the differences between “rich” and “specific” prompts, which serve as the basis for the AI model’s generation process and directs the technology’s output).

⁵⁰ See *Mr Li v Miss Liu*, Beijing Internet Ct., Case No. 11279, Nov. 27, 2023 (China) (recognizing copyrightability of images created by AI where a human input prompt words and parameters).

⁵¹ Charles Ross, *Does ChatGPT Give the Same Answer to Everyone?*, MEDIUM (Mar. 20, 2023), <https://medium.com/@charles-ross/does-chatgpt-give-the-same-answer-to-everyone-521e3e9355a4>.

⁵² *Id.*

9

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

applications are thus labeled as “black boxes” because their decision-making processes are inaccessible and uninterpretable by human beings.⁵³

For example, consider a situation involving Midjourney, a generative AI program that produces images from natural human language descriptions.⁵⁴ Two artists simultaneously prompt Midjourney: one writes in colloquial English and the other writes in colloquial French.⁵⁵ Midjourney will deliver different results.⁵⁶ The artists themselves did not make creative choices that differ from one another simply by communicating in their native languages; rather, it is the AI models’ internal methodology and training that results in different artistic outputs.⁵⁷ Even so, the two artists may eventually learn how to prompt the Midjourney tool to more precisely reflect their respective visions.⁵⁸ With current technology, the problem remains that the algorithm is ultimately remixing the content upon which it has been trained, making a series of binary choices, and providing an output (thereby doing the “creating”), while the human is simply inputting a series of word prompts.⁵⁹

⁵³ A. Bandi, P.V.S.R. Adapa & Y.E.V.P.K. Kuchi, *The Power of Generative AI: A Review of Requirements, Models, Input–Output Formats, Evaluation Metrics, and Challenges*, 15 FUTURE INTERNET at 49 (2023).

⁵⁴ *About*, MIDJOURNEY, midjourney.com/home (last visited Feb. 28, 2024); Richard Greenan, *Midjourney AI: The Complete Guide to the AI Art Generator*, CAREERFOUNDRY (Dec. 21, 2023), <https://careerfoundry.com/en/blog/ux-design/midjourney/>.

⁵⁵ See Allison Snyder, *AI’s Language Gap*, AXIOS (Sept. 8, 2023), <https://www.axios.com/2023/09/08/ai-language-gap-chatgpt> (describing AI’s current English bias across disciplines, as most of today’s most popular generative AI tools are built and trained on texts and data in English and Chinese).

⁵⁶ *Id.* (stating that AI models can translate prompts and responses into English well, but often makes mistakes when translating English into other languages).

⁵⁷ See, e.g., *ChatGPT Prompt for Bilingual Language*, OPENAI, <https://community.openai.com/t/chatgpt-prompt-for-bilingual-language/325703> (last visited May 14, 2024) (providing examples on how prompts in English yield different results from prompts in Vietnamese).

⁵⁸ See Mark C. Humphrey, *Not-So-Starry Night for Copyright Applicant: Copyright Office Again Refuses to Register an AI-Generated Image*, MSK BLOG (Dec. 12, 2023), <https://www.msk.com/newsroom-alerts-copyright-office-ai-registration-refusal> (noting the U.S. Copyright Office’s Review Board’s rejection for reconsideration on a refusal to register an image created with generative AI for failure to meet the human authorship requirement, emphasizing that the Office expects applicants will focus on arguing to their own creative contributions to their AI inputs). The Review Board stated that “creative contributions” to AI prompts include selection, arrangement, complexity, and creativity of material used in input prompts, highlighting the uncertainty as to whether the US Copyright Office will be open to these arguments despite hewing closely to established precedent. *Id.*

⁵⁹ See *id.* (describing the Copyright’s Office reasoning in denying copyright registration to an AI that remixed the content it was trained on to produce an output because the creation was a function of how the model work and the training images rather than the choices made by the human).

2. The Beijing Internet Court Decision on Copyrighting AI Generated Images

In late November 2023, the Beijing Internet Court in China recognized copyright in AI-generated images.⁶⁰ The decision focused on an individual's use of the generative AI product Stable Diffusion to generate an image of a fictional woman.⁶¹ A blogger reused the image in his own blog post, spurring litigation.⁶² The Internet Court found that the image, although generated by artificial intelligence, met the requirements of originality required by Chinese law because the photo reflected a human being's original intellectual investment.⁶³ In conceiving the image, designing the presentation of characters, choosing prompt words, arranging the order of the prompt words, setting relevant parameters, and selecting the final iteration, the Internet Court determined the originality standard was satisfied.⁶⁴ The images involved in the case met the threshold requirements of "intellectual achievements," thereby earning copyright protection.⁶⁵ This decision contrasts starkly with an earlier ruling by the same court in 2019.⁶⁶

The generative AI user in this case selected a specific manifestation of Stable Diffusion, with some predefined parameter settings, as well as the different models to be used and a positive prompt (what one wants to see in the generated image).⁶⁷ He added a negative prompt (what one does not want to see), and then selected a number of thirty-three iterations to generate the image.⁶⁸ He also changed the aspect ratio to the standard 2:3 proportions, modified a parameter that describes how closely the model follows the prompt, and selected a seed number.⁶⁹ These selections led to the final creation, which was a product of intellectual labor, originality, and perspective.⁷⁰ Although the Internet Court characterized these activities as the plaintiff's active participation and own artistic intervention in creating the work, it is argued that the work did not actually result from identifiable

⁶⁰ See *Mr Li v Miss Liu*, Beijing Internet Ct., Case No. 11279, Nov. 27, 2023 (China) (holding that an AI-generated image was entitled to copyright protection).

⁶¹ *Id.*

⁶² *Id.* at 1–2.

⁶³ *Id.* at 10–14.

⁶⁴ *Id.*

⁶⁵ Aaron Wininger, *Beijing Internet Court Recognizes Copyright in AI-Generated Images*, NAT'L L. REV. (Nov. 29, 2023), <https://www.natlawreview.com/article/beijing-internet-court-recognizes-copyright-ai-generated-images>. The Beijing Internet Court defined "intellectual achievements" as the results of intellectual activities such that the work reflects the intellectual input of a natural person. *Mr Li v Miss Liu*, Beijing Internet Ct., Case No. 11279, Nov. 27, 2023, at 10–14 (China).

⁶⁶ *Beijing Feilin v. Baidu*, Beijing Internet Ct., Case No. 239, Apr. 26, 2019 (China) (holding that, although an AI-generated report satisfied the originality requirement, the report failed to meet the natural person authorship requirement because it was "created" by AI software).

⁶⁷ *Mr Li v Miss Liu*, Beijing Internet Ct., Case No. 11279, Nov. 27, 2023, at 7–8 (China).

⁶⁸ *Id.* at 4–5.

⁶⁹ *Id.* at 6, 8; see Jakub Wyczik, *Threshold of Originality for AI Output in Copyright Law*, LINKEDIN (Dec. 6, 2023), <https://www.linkedin.com/pulse/threshold-originality-ai-output-copyright-law-jakub-wyczik-hh9sf> (providing a technological interpretation of the Internet Court's description of the defendant's actions and the plaintiff's intellectual contributions to the resulting work).

⁷⁰ *Mr Li v Miss Liu*, Beijing Internet Ct., Case No. 11279, Nov. 27, 2023, at 10–14 (China).

11

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

human creativity due to the nature and mechanics of generative AI tools.⁷¹ In other words, while setting algorithmic parameters and inserting a variety of words takes minimal time, effort, and ideation, the actual output of Stable Diffusion could not be said to reflect the user’s artistic creativity.⁷² This is largely due to the fact that a variety of elements in the generation process are random and depend on aspects of the process that are wholly uncreative.⁷³

The Beijing Internet Court’s decision is at odds with existing U.S. and international copyright law jurisprudence in several ways.⁷⁴ First and foremost, it conflicts with the established idea that AI-generated works are, by dint of their being products of generative AI, remixes of the copyrighted works on which they are trained without permission.⁷⁵ As explored above, the “remixes,” while directed by human prompting, do not amount to human creativity as it has been traditionally conceived and, depending on the hardware, the algorithm, and the data set, the same set of prompts will achieve different results that negate any direct tie between the prompt and the result.⁷⁶

3. Copyright in AI Generated Images in the United States

In the United States, the question of whether the originality in AI prompts is sufficient to overcome the creativity requirement for copyright registration has not yet been directly addressed by the U.S. Copyright Office.⁷⁷ In a recent and well-known application for copyright registration, an individual who submitted a work

⁷¹ See Olivia Rafferty, *Chinese Copyright Ruling on AI-Generated Images Leads to Fallout Overseas*, WORLD TRADEMARK REV. (Dec. 15, 2023), <https://www.worldtrademarkreview.com/article/chinese-copyright-ruling-ai-generated-images-leads-fallout-overseas> (citing arguments that AI should be considered a replacement for the human author because the resulting works are autonomously completed by AI); Tingting Wen, *Beijing Internet Court Recognizes Copyright in AI-Generated Image*, 19 J. OF INTELL. PROP. L. & PRAC. 203 (2024), <https://doi.org/10.1093/jiplp/jpad127> (discussing the intricacies of the Internet Court’s decision).

⁷² See Rafferty, *supra* note 71 (presenting arguments from Chinese attorneys and legal scholars that adjustment of parameters and selection of keywords within an AI prompt should not constitute any type of human creativity as it is traditionally understood).

⁷³ See Wen, *supra* note 71 (describing how the generation process of utilizing the AI in question involved somewhat random elements and depended on variables such as the equipment available).

⁷⁴ Wyczik, *supra* note 69.

⁷⁵ See Mike Taylor, *Using AI Without Exploiting Artists*, LADDER (Jan. 26, 2024), <https://ladder.io/blog/using-ai-without-exploiting-artists> (explaining that generative AI models like Stable Diffusion were trained on millions of images scraped from the internet without the copyright holder’s permission, which artists claim violates copyright law).

⁷⁶ Snyder, *supra* note 55; see *Thaler*, 2023 U.S. Dist. LEXIS 145823, at *18–20 (holding that plaintiff’s inability to identify a court decision recognizing copyright in a work that originated with a non-human and the attenuation of human creativity from the actual generation of the final work requires a finding that copyright cannot be granted to AI-generated works despite the human’s inputs through prompts).

⁷⁷ See Caldwell, *supra* note 17, at 436–42 (explaining current US jurisprudence, or the lack thereof, regarding creativity in AI prompts for the purposes of copyright protection through a historical perspective).

did not list *himself* as an author but rather the AI; the Office thus denied the registration.⁷⁸ Had he claimed authorship himself, his choices, even if numerous, were still likely not the kinds of choices that would result in an expected outcome or benefit from free and unbound human creativity.⁷⁹

Furthermore, on December 11, 2023, the U.S. Copyright Office once again issued a letter refusing to register a work created by generative AI.⁸⁰ The work in question is a photograph by the putative registrant enhanced by a generative AI painting application.⁸¹ The author of the photograph applied a “style” input from Vincent Van Gogh’s *The Starry Night* to the generative AI app and chose how “much” of the style to transfer.⁸² The Copyright Office logically found that the AI application’s interpretation of the applicant’s photograph in the style of another painting was simply a function of how the model worked together with the images on which it was trained, which was ultimately not attributable to specific contributions or instructions received from the applicant.⁸³ Although the applicant selected the numerical variable for the potency of the style, that choice alone was insufficient to qualify for copyright protection; that kind of selection belongs to the category of *de minimis* authorship not protected by copyright.⁸⁴

4. Analogous Technologies and Creativity

Debates and litigation surrounding authorship when technology is utilized in the creation of a work or when an author partners with another individual to carry out their vision are long-standing and contentious.⁸⁵ New York society photographer Napoleon Sarony took a photograph of Oscar Wilde in 1882, when

⁷⁸ See Second Request for Reconsideration for Refusal to Register A Recent Entrance to Paradise, SR # 1-7100387071, Correspondence ID: 1-3ZPC6C3 at 3 (U.S. Copyright Off. Feb. 14, 2022) (second admin. review), <https://www.copyright.gov/rulings-filings/review-board/docs/a-recent-entrance-to-paradise.pdf>.

⁷⁹ See *id.* (stating that the applicant must either provide evidence the work is the product of sufficient human authorship, which would be a difficult feat given current understandings and definitions of “original work of authorship,” or convince the U.S. Copyright Office to depart from long-standing copyright jurisprudence and precedent); Wen, *supra* note 71 (explaining that using AI to generate works involves random elements and depends on non-creative variables such as the equipment used).

⁸⁰ Second Request for Reconsideration for Refusal to Register SURYAST, SR # 1-11016599571, Correspondence ID: 1-5PR2XKJ at 9 (U.S. Copyright Off. Dec. 11, 2023) (second admin. review), <https://www.copyright.gov/rulings-filings/review-board/docs/SURYAST.pdf>.

⁸¹ *Id.* at 2–3.

⁸² *Id.*

⁸³ *Id.* at 8.

⁸⁴ *Id.*

⁸⁵ See, e.g., James Grimmelman, *There’s No Such Thing as a Computer-Authored Work – And It’s a Good Thing, Too*, 39 COLUM. J. L. & ARTS. 403, 406 (2016) (arguing that every category of copyrightable work can be and regularly is created using computer technology); Nahide Basri, *The Question of Authorship in Computer-Generated Work*, UNIV. OF PENN. J. OF L. & SOC. CHANGE (Jan. 12, 2020), <https://www.law.upenn.edu/live/news/9691-the-question-of-authorship-in-computer-generated> (discussing the debate surrounding computer-generated works and presenting arguments for the creation of computer-specific copyright law frameworks).

13

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

photography was a relatively new medium, and claimed copyright in that photograph.⁸⁶ The Supreme Court confirmed, in 1884, that this photograph was copyrightable, not because the act of photographing confers authorship, but because of the selections and arrangements Sarony made in posing his subject, positioning the drapes and lighting, suggesting and evoking the desired expression, and otherwise creatively composing the eventual image.⁸⁷ This general principle has been reaffirmed and built upon in case law until present day.⁸⁸ It is not the snapping of the camera button that confers authorship; it is the indicia of human creativity apparent in the final photograph that does so.⁸⁹ Under this standard, not every photograph is copyrightable.⁹⁰ The nature of artificial intelligence calls for a renewed focus on the salient creative actions an author takes which confer authorship on him and not on an agent or on technology.⁹¹

II. THE CREATIVITY REQUIREMENT IN RELATION TO GENERATIVE AI

Having provided vital context of the recent findings and controversies in the intersection of copyright law and generative AI, this Part discusses copyright law's requirement for originality and creativity and applies it to generative AI.⁹² Section A examines existing law and relevant legal tests, Section B considers human creativity, or the lack thereof, in prompting generative AI, and Section C underscores the importance of copyright's elasticity.⁹³ Ultimately, Part II concludes that generative AI does all the conception and execution of potentially copyrightable works on its own because the technology suggests random collages

⁸⁶ *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 54 (1884).

⁸⁷ *Id.* It is important to note that Oscar Wilde and Napoleon Sarony contractually concluded between themselves that any authorship would belong to Sarony, and thus the Court did not address this issue. *Id.*

⁸⁸ See, e.g., *SHL Imaging, Inc. v. Artisan House, Inc.*, 117 F. Supp. 2d 301, 310 (S.D.N.Y. 2000) (describing how the copyrightability of a photograph boils down to the photographer's creative control thereof); *Schrock v. Learning Curve Int'l, Inc.*, 586 F.3d 513, 519 (7th Cir. 2009) (indicating that the creative control of the photographer informs the copyrightability of the work).

⁸⁹ *Burrow-Giles*, 111 U.S. at 60 (holding that photographic works were protected by copyright, as the photos could be traced back to deliberate and purposeful choices of the photographer such as lighting and positioning).

⁹⁰ See *id.* (describing some of the creative activities a photographer undertakes to impart copyrightable originality in the resultant photograph, including selecting costumes, draperies, positioning, lighting, expressions, dispositions, and arrangements, while holding that copyright applicants must meet this threshold requirement of creativity and authorial contribution).

⁹¹ See Ginsburg & Budiardjo, *supra* note 19, at 360 (describing how a principal author controls the assistant when the principal influences what the assistant does and how they accomplish the task); Roosa Wingström, Johanna Hautala & Riina Lundman, *Redefining Creativity in the Era of AI? Perspectives of Computer Scientists and New Media Artists*, 36 CREATIVITY RSCH. J. 177, 186–88 (2014), <https://www.tandfonline.com/doi/full/10.1080/10400419.2022.2107850> (discussing the various elements of creativity and the ways in which AI impacts and challenges the traditional understanding of independent creativity).

⁹² See *infra* notes 96–207 and accompanying text.

⁹³ See *infra* notes 96–207 and accompanying text.

of the training data.⁹⁴ This Part then argues that no copyrightable human creativity, absent the works on which the AI is trained, is performed in the production of the images or text the AI generates, which negates the necessary creative human contribution required for copyright protection.⁹⁵

A. Generative AI Existing Law and Relevant Tests

Courts and academics have respectively proposed and employed standards to determine whether a particular use of generative AI involves enough human creativity to confer the status of authorship on a generative AI user for a particular work.⁹⁶ While definitions of “generative AI” vary, the current definition of an “AI system” provided by the Organisation for Economic Co-operation and Development (OECD) creates a useful starting point.⁹⁷ An AI system is a machine-based system that infers, from the inputs given to it, how to generate outputs such as predictions, recommendations, content, or decisions that can influence physical or virtual environments based on various levels of autonomy and adaptiveness utilized after deployment.⁹⁸ Using this definition as a baseline description of what generative AI is and how it works, the following subsections explore how current copyright jurisprudence could interpret the relationship between generative AI and the types of undertakings legally required for authoring a copyrightable work.⁹⁹

1. Proposed Causation Tests for Generative AI

Professor Daniel Gervais recently tackled the authorship question in relation to generative AI.¹⁰⁰ He proposes using “originality causation” as the test to determine whether a human has contributed the requisite amount of creativity to a work such that any AI technology involved could be considered a tool, as opposed to the rightful author of the work.¹⁰¹ Professor Gervais posits that to apply originality causation, courts should identify all machine-made choices in the creation process and exclude them in determining whether a work is original for purposes of copyright law.¹⁰² In other words, he believes that if all or a significantly large portion of relevant choices were made by a machine, the putative work is not

⁹⁴ See *infra* notes 160–181 and accompanying text.

⁹⁵ See *infra* notes 182–186 and accompanying text.

⁹⁶ See, e.g., Edward Lee, *Prompting Progress: Authorship in the Age of AI*, 76 FLORIDA L. REV. at 18 (Oct. 22, 2023) (forthcoming), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4609687 (exemplifying an academic discussion and historical analysis of the evolution of copyright law standards in the context of the debate surrounding AI and authorship).

⁹⁷ *Recommendation of the Council on Artificial Intelligence*, OECD (May 22, 2019), <https://oecd.ai/en/assets/files/OECD-LEGAL-0449-en.pdf>.

⁹⁸ *Id.*

⁹⁹ See *infra* notes 100–158 and accompanying text.

¹⁰⁰ Daniel Gervais, *The Machine as Author*, 105 IOWA L. REV. 2053, 2098–106 (2020) (examining the threshold of autonomy that separates the output of AI from the humans that programmed the models or used the technology).

¹⁰¹ *Id.* at 2098–101.

¹⁰² *Id.* at 2100–01.

15

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

protected by copyright; and inversely, that if the majority of the relevant choices were made by a human, the putative work has the capacity for copyright protection.¹⁰³ He argues that causation is a well-known phenomenon in other areas of law like product liability, and that, in the context of AI and copyright, the originality causation test enables the separation of protectable human creative expression from non-protectable machine outputs.¹⁰⁴

Professor Gervais’s causation analysis follows Professor Dan Burk’s proposal two years earlier that authorship can best be evaluated and understood in terms of causation, intent, and volition, amongst other related legal doctrines.¹⁰⁵ With respect to machine intelligence specifically, Professor Burk notes that these technologies do not behave voluntarily, purposefully, or autonomously.¹⁰⁶ In comparing the behavior of machine learning systems with that of paint-splatter artist Jackson Pollock, Professor Burk highlights that not every Pollock-like painting will result in the requisite level of creativity and therefore will not deserve copyright protection for want of an author.¹⁰⁷ For example, if Pollock were to begin setting up a paint splatter machine and gets interrupted, before he could finish his set-up, by the machine short-circuiting and spreading paint around the studio, Pollock could not be considered the author of the result of these circumstances even if the machine struck a canvas with paint.¹⁰⁸ Put in a more technology-forward way, “[w]hile some element of randomness does not eliminate authorship, the putative author must be able to constrain or channel the program’s processing of the source material.”¹⁰⁹ In other words, copyright law does not expect a supposed author to be in complete and domineering control over every aspect of the work they create; there is room for experimentation and error and serendipity.¹¹⁰ But an author must still be the source—or the “cause”—of the traditional elements of authorship.¹¹¹ When a person prompts an AI machine to produce a drawing of the Statue of Liberty viewed from the southern tip of Manhattan, the human is acting as the commissioner of the resultant image by issuing instructions and the AI technology is acting as the artist.¹¹² Someone who commissions art may eventually be the right

¹⁰³ *Id.*

¹⁰⁴ *Id.* at 2099–101.

¹⁰⁵ Dan L. Burk, *Thirty-Six Views of Copyright Authorship, by Jackson Pollock*, 58 HOUS. L. REV. 263, 321 (2020).

¹⁰⁶ *Id.* at 319–20.

¹⁰⁷ *See id.* at 298–300 (relying on the legal doctrine of causation to argue that when an unforeseen intervening cause breaks the causal chain between the artist and the resulting work is both unforeseeable and not a direct result of Pollock’s original expression).

¹⁰⁸ *Id.*

¹⁰⁹ Kernochan Center for Law, Media and the Arts, Columbia Law School, Comment Letter on Notice of Inquiry on Artificial Intelligence and Copyright, at 5 (Oct. 30, 2023), <https://www.regulations.gov/comment/COLC-2023-0006-9090> [hereinafter Kernochan Comment].

¹¹⁰ *See* Shyamkrishna Balganes, *Causing Copyright*, 117 COLUM. L. REV. 1, 66 (stating that an author’s complete control over the resulting work is not necessary to satisfy the legal-creation test).

¹¹¹ U.S. COPYRIGHT OFF., L.C. NO. 10-35017, SIXTY-NINTH ANNUAL REPORT OF THE REGISTER OF COPYRIGHTS 5 (1967).

¹¹² *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, 88 Fed. Reg. 16190 (Mar. 16, 2023) (to be codified at 37 C.F.R. pt. 202).

holder in the work, either by way of the work made for hire doctrine or contractual agreement, but they are not the author or artist in whom original copyright inheres.¹¹³

2. Causation Tests Application

Like many areas of copyright law, applying causation tests to determine authorship requires fact-specific inquiries into how a work came into being.¹¹⁴ In another analog context, one might question whether an aspiring poet who writes a poem comprised solely of every eleventh word of the next *New York Times* article they read could constitute a copyrightable work.¹¹⁵ In the terminology articulated by Professors Gervais and Burk, is the resultant poem *caused* by the poet's *originality*?¹¹⁶ Are they an author of a copyrightable work?¹¹⁷ Reasonable minds might differ on the answer.¹¹⁸ One might argue that authorship is perhaps based on the arrangement of the words: for example, did the poet put four words in every line to make the poem flow with some kind of sonorous advantage when spoken aloud?¹¹⁹ Did they select a specific *New York Times* article that would be

¹¹³ U.S. COPYRIGHT OFFICE, CIRCULAR 30: WORKS MADE FOR HIRE 1, 4 (2021), <https://www.copyright.gov/circs/circ30.pdf>.

¹¹⁴ See Burk, *supra* note 105, at 267 (stating that copyright, like other areas of law where liability is assigned on the basis of causality, requires factual inquiry into the many factors that contribute to the expression of a particular work in order to determine a point of origin where authorship can attach).

¹¹⁵ See U.S. COPYRIGHT OFFICE, CIRCULAR 14: COPYRIGHT IN DERIVATIVE WORKS AND COMPILATIONS 1–2 (2020), <https://www.copyright.gov/circs/circ14.pdf> [hereinafter CIRCULAR 14: DERIVATIVE WORKS] (providing guidance on the copyrightability of a derivative work, including the requirement of incorporating some of a preexisting work along with the addition of new original copyrightable authorship). The author used a recent *New York Times* article to conceptualize what a poem comprised of every eleventh word would look like: Cuts grand held the continued names shifts the 1940 soon another and half Amsterdam's address mapping archival a Amsterdam that millimeter birdsong some English-language actor the it without drops the hopscotches instance what garden metallic voices and an a suicide Germany recounting could man's employee days the linkages cozy you where McQueen protesting images with into history Jewish a lived from baby occupied that most in. See Manohla Dargis, 'Occupied City' Review: Mapping the Holocaust, Street by Street, N.Y. TIMES (Dec. 24, 2023), <https://www.nytimes.com/2023/12/24/movies/occupied-city-review-mapping-the-holocaust-street-by-street.html> (providing an article for the purposes of the thought experiment above).

¹¹⁶ See Gervais, *supra* note 100, at 2100 (framing the originality test using a causation framework); Burk, *supra* note 105, at 265–67 (arguing that causal tracing is a useful framework for identifying and assigning authorship of works for purposes of copyright).

¹¹⁷ See Gervais, *supra* note 100, at 2073–85 (discussing the different ways authorship can be achieved for purposes of a copyrighted work through a historical analysis); Burk, *supra* note 105, at 265 (raising the question as to whether works computer systems and AI technology could or should be considered authors for copyright purposes).

¹¹⁸ See Gervais, *supra* note 100, at 2088–98 (noting the different perspectives as to what deems authorship for copyright purposes).

¹¹⁹ See, e.g., U.S. COPYRIGHT OFFICE, CIRCULAR 33: WORKS NOT PROTECTED BY COPYRIGHT 3 (2021), <https://www.copyright.gov/circs/circ33.pdf> (explaining that copyright may be available for a purposeful arrangement of specific content deliberately selected and arranged in a creative way).

17

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

particularly well-suited to this hunt-and-peck method?¹²⁰ In contrast, one could argue that the poet is not an author, especially if the poem comprises a single line of words arranged in the same order as they appeared in article they were taken from.¹²¹ Under this claim, the resultant “poem” is merely an unprotectable “style” of writing that selects every eleventh word, and not an actual creative work.¹²²

3. The U.S. Copyright Office’s Current Test for Creativity in Originality

Familiar symbols and designs, or a simple combination of a few familiar symbols or designs, are uncopyrightable and cannot be registered with the U.S. Copyright Office.¹²³ Conversely, a work of authorship that *incorporates* one or more familiar symbols or designs into a larger design may be registered if the work as a whole contains a sufficient amount of creative expression.¹²⁴ In order for copyright protection to subsist, the work must have been independently created and demonstrate a modicum of creativity.¹²⁵ There are many examples illustrating how the U.S. Copyright Office has grappled with questions regarding creativity thresholds that echo the aforementioned debates surrounding generative AI.¹²⁶

For example, various arrangements of simple shapes may or may not be copyrightable; it would depend on whether the arrangements demonstrate some minimal creativity.¹²⁷ The U.S. Copyright Office denied copyright protection to two graphic logos proposed by Crowne Melbourne and Cerner Corporation—one in the shape of a crown achieved by placing variously sized dots together and another in the shape of the letter C with two parallel lines swooshing through it, respectively—because the Office could not identify the requisite amount of

¹²⁰ See *id.* (explaining that the copyright protection would be limited to the selection and arrangement of the specific content, not that of any content in a particular way).

¹²¹ See CIRCULAR 14: DERIVATIVE WORKS, *supra* note 115, at 1 (stating that a derivative work must be sufficiently different from the work on which it was based to be considered a “new work” and thus be copyrightable, as minor changes of little substance does not constitute a new version).

¹²² See *A Guide to Copyright Infringement*, THE ILLUSTRATOR’S GUIDE (Apr. 5, 2024), <https://theillustratorsguide.com/copyright-infringement/> (explaining that “style” cannot be copyrighted, and creators can make their own works in a similar style without infringing on copyright).

¹²³ COMPENDIUM OF COPYRIGHT PRACTICES, *supra* note 15, at § 906.2.

¹²⁴ *Id.*

¹²⁵ *Feist Publications, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340 (1991).

¹²⁶ Letter from Shira Perlmutter, Register of Copyrights and Director, U.S. Copyright Office, to Chis Coons, Chair, Subcommittee on Intellectual Property of the U.S. Senate (Feb. 23, 2024), <https://copyright.gov/laws/hearings/USCO-Letter-on-AI-and-Copyright-Initiative-Update-Feb-23-2024.pdf?locl=blogcop> (presenting to Congress the steps and measures the U.S. Copyright Office took in 2023, and plans to take in 2024, to examine the implications of the existing copyright requirements in the context of generative AI).

¹²⁷ See COMPENDIUM OF COPYRIGHT PRACTICES, *supra* note 15, at § 906.2 (highlighting how some combinations of common geometric shapes may contain sufficient creativity with respect to how they are juxtaposed or arranged to support a copyright, but not every combination or arrangement will be sufficient to meet this test).

creativity in each logo.¹²⁸ In the case of the Crown Melbourne image (exhibited below), the Office found the work consisted of a single repeated element, solid circles, arranged in symmetrical rows; and while the rows were curved, there was insufficient creativity for protection.¹²⁹ Furthermore, the Office held that the idea of a five-point crown was familiar and unprotectable as well.¹³⁰ Similarly, the proposed logo of the Cerner Corporation was denied protection.¹³¹ Neither the design, the color, or the combination thereof were deemed original enough to qualify for copyright protection.¹³²

The refusal of copyright registration to these two designs illustrates the creativity hurdle that must be surpassed to achieve copyrightability in the United States.¹³³ A “modicum of creativity” is meaningless without context and examples, yet the particular designs demonstrate that “modicum” indeed means something more than nothing; it is more than a crown fashioned out of circles and more than a multicolored “C” with lines through its center.¹³⁴ There was no known generative AI involved in the creation of either logo, which provides a useful comparison for then adding in the technological AI overlay to the analysis.¹³⁵ As one commenter has stated, the test for originality does not lie so much in the question of whether a generative AI tool was used or not, but rather in the question of whether there has

¹²⁸ See Second Request for Reconsideration for Refusal to Register Crown Design, SR # 1-10320239801, Correspondence ID: 1-50LFG3Z (U.S. Copyright Off. Jan. 27, 2023) (second admin. review), <https://copyright.gov/rulings-filings/review-board/docs/Crown-Design.pdf> (explaining that the repetition of a single element—common circles—as selected and arranged in the application fails to evince sufficient creativity to satisfy the threshold creativity requirements); Second Request for Reconsideration for Refusal to Register C Design, SR # 1-7053886421, Correspondence ID: 1-3VN8JZW (U.S. Copyright Off. Oct. 17, 2002) (second admin. review), <https://copyright.gov/rulings-filings/review-board/docs/c-design.pdf> (explaining the decision that a logo consisting of a letter C with two streaks extending horizontally and parallel to each other from the middle of the letter C is not copyrightable). See *infra* notes 138–139 and accompanying text (displaying the proposed designs and logos).

¹²⁹ See Second Request for Reconsideration for Refusal to Register Crown Design, *supra* note 128, at 2–3 (describing how neither the individual elements nor their combination constitutes sufficient creativity for copyright protection).

¹³⁰ *Id.* at 3.

¹³¹ Second Request for Reconsideration for Refusal to Register C Design, *supra* note 128, at 6.

¹³² *Id.* at 2.

¹³³ See *id.* at 3–5 (reiterating the legal creativity requirements a work must satisfy in order to obtain copyright protection and exemplifying a denial in copyright registration for a design with too few creative elements); Second Request for Reconsideration for Refusal to Register Crown Design, *supra* note 128, at 2–3 (serving as another example of copyright registration denial for a logo with minimal design features lacking sufficient creativity).

¹³⁴ See *Feist Publ'ns, Inc., v. Rural Telephone Serv. Co.*, 499 U.S. 340, 358 (1991) (noting that although a modicum of creativity is a relatively small hurdle, a work of authorship still needs at least some creativity to be eligible for copyright protections); Second Request for Reconsideration for Refusal to Register C Design, *supra* note 128, at 3 (stating that although only a modicum of creativity is necessary, even some applicants fail to meet this low threshold – such as the stylized blue letter C with the two parallel green streaks in the application).

¹³⁵ See Second Request for Reconsideration for Refusal to Register C Design, *supra* note 128, at 2–3 (explaining how the design was created with no mention of the use of generative AI); Second Request for Reconsideration for Refusal to Register Crown Design, *supra* note 128, at 2–3 (detailing the proposed design and logo without reference to the use of generative AI in the creation process).

19

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

been sufficient or insufficient human creative contribution.¹³⁶ The “test” then is as simple, or perhaps as complex, as determining what the correct amount of human creative contribution is to result in the requisite originality.¹³⁷



Crown Melbourne Limited's Proposed Logo¹³⁸



Cerner Corporation's Proposed Logo¹³⁹

Registrations for two non-AI-assisted designs refused by the U.S. Copyright Office.

4. Comparing Analog and Generative AI Environments

The question of creative decision-making in the analog environment is not so different from that in the AI environment.¹⁴⁰ If a putative author's decision is driven by too many external considerations lacking creative freedom, it cannot be considered creative for purposes of copyright protection.¹⁴¹ This has been true at

¹³⁶ See Kateryna Milityna, *Human Creative Contribution to AI-Based Output – One Just Can't Get Enough*, 72 GRUR INT'L 939, 941–43 (2023) (attempting to clarify confusion of copyrightability of AI-generated works by presenting a five-part test for distinguishing sufficient human creative participation in AI-based outputs for copyright purposes).

¹³⁷ See *id.* (highlighting the vast variety of circumstances that could complicate or inform the seemingly simple creativity threshold inquiry).

¹³⁸ Second Request for Reconsideration for Refusal to Register Crown Design, *supra* note 128, at 1.

¹³⁹ Second Request for Reconsideration for Refusal to Register C Design, *supra* note 128, at 1.

¹⁴⁰ See Milityna, *supra* note 139, at 946–48 (analyzing and analogizing the user's choices and the AI technology's choices through the creation process by identifying the choices and classifying them as steps requiring creativity or steps completely devoid of creative freedom).

¹⁴¹ *Id.* at 943 n.62; see *Trends and Developments in Artificial Intelligence: Challenges to the Intellectual Property Rights Framework*, at 70–74 (2020), <https://digital-strategy.ec.europa.eu/en/library/trends-and-developments-artificial-intelligence-challenges->

least since a sinuous and aesthetically award-winning bicycle rack was denied copyright protection by the U.S. Court of Appeals for the Second Circuit in 1987.¹⁴² In reaching this decision, the court noted that the choices were dictated by preexisting physical constraints of temporary bicycle storage rather than the creativity of a putative author; the principle has since been oft-repeated in case law over ensuing decades.¹⁴³

For a putative author working within a constrained environment, there is a higher hurdle to clear to demonstrate that one's choices were creative as opposed to predetermined or dictated by practicalities.¹⁴⁴ This concept applies to photography as well.¹⁴⁵ For example, consider all of the photographs of the Statue of Liberty that have been taken from the southern tip of Manhattan at sunset.¹⁴⁶ Copyright in any given photograph will comprise the creative elements it displays, and not the subject itself or any of the qualities that exist in other similar photographs—qualities considered to be outside the creative control of the photographer and thus not copyrightable.¹⁴⁷

This analysis parallels the analysis for copyright using the tool of generative AI.¹⁴⁸ In other words, to the extent human creativity can be ascertained in a

intellectual-property-rights-framework (summarizing copyright jurisprudence implying that if the author makes only minimal or non-creative choices, copyright protection may not be afforded for failure to satisfy the creative threshold through various and personal choices).

¹⁴² See *Brandir Int'l, Inc. v. Cascade Pac. Lumber Co.*, 834 F.2d 1142, 1147–49 (2d Cir. 1987) (holding that a bicycle rack is not copyrightable because the form was too significantly influenced by utilitarian concerns and lacked sufficient creative freedom).

¹⁴³ See *id.* at 1147 (explaining that the bicycle rack is a product of industrial design as form and function are inextricably connected to the rack, and the resulting design is the result of utilitarian necessities rather than creative choices).

¹⁴⁴ See, e.g., Second Request for Reconsideration for Refusal to Register “nazStreetwise.mib”, SR # 1-5900194675, Correspondence ID: 1-31ZTZRS at 6–8 (U.S. Copyright Off. Sept. 27, 2019) (second admin. review) (explaining common law jurisprudence that limits copyright protection for elements of a work that are dictated by external factors, mechanical specifications, necessity, utility, and other demands, ultimately denying this applicant due to lack of creative choices available to the author despite explanation of those choices by applicant).

¹⁴⁵ See *Leigh v. Warner Bros., Inc.*, 212 F.3d 1210, 1215 (11th Cir. 2000) (holding that any individual photograph of the so-called “Bird Girl” statue in a Georgia cemetery is only copyrightable insofar as its creative aspects, and that any “mood” or “feel” evoked by the statute, location, or overall combination of the subject matter is non-original and unprotectable).

¹⁴⁶ See *Statue of Liberty*, GETTY IMAGES, <https://www.gettyimages.com/photos/statue-of-liberty> (last visited May 15, 2024) (providing a database of nearly 24,000 photos of the Statue of Liberty from various angles and vantage points); Suvid Kakkar, Photograph of the Statue of Liberty (July 18, 2017), https://upload.wikimedia.org/wikipedia/commons/6/69/Statue_of_Liberty_Silhouette.jpg (displaying an image of the Statue of Liberty).

¹⁴⁷ COMPENDIUM OF COPYRIGHT PRACTICES, *supra* note 15, at §§ 909–909.2; see *Leigh v. Warner Bros.*, 212 F.3d. at 1214–15 (holding that the photographer’s copyright does not cover the subject matter itself or the surrounding setting, upholding longstanding jurisprudence that artists do not possess copyright in the “reality of [their] subject matter”).

¹⁴⁸ See Martin Gomez, *Is It Possible to Copyright Works that Include AI-Generated Material?*, GOODWIN (Oct. 2, 2023), <https://www.goodwinlaw.com/en/insights/publications/2023/10/insights-technology-aiml-is-it-possible-to-copyright-works> (stating that U.S. law has long recognized the use

21

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

photograph of the Statue of Liberty—or the output of generative AI—there is at least the possibility of copyrightable authorship.¹⁴⁹



Suvid Kakkar, Photograph of the Statue of Liberty, 2017, CC-BY.

5. Challenges to the U.S. Copyright Guidance on Registration for Technology

The most recent version of the U.S. Copyright Office Compendium details the requisite authorship elements as they were described in 1966. Importantly, the language states:

[T]he Office will not register works produced by a machine or mere mechanical process that operates randomly or automatically without any creative input or intervention from a human author. The crucial question is “whether the ‘work’ is basically one of human authorship, with the computer [or other device] merely being an assisting instrument, or whether the traditional elements of authorship in the work (literary, artistic, or musical expression or elements of selection, arrangement, etc.) were actually conceived and executed not by man but by a machine.”¹⁵⁰

Professor Edward Lee, with a forthcoming article in *The Florida Law Review*, recently argued that the existing U.S. Copyright Office guidance on

of technology as one component of a human-led creative outcome and arguing that AI-generated models simply act as tools that support and augment the human being’s creativity process).

¹⁴⁹ See Request for Reconsideration for Refusal to Register *Zarya of the Dawn*, SR # VAu001480196, Correspondence ID: 1-5GB561K at 12 (U.S. Copyright Off. Feb. 21, 2023) (first admin. review) (granting a narrowly tailored copyright registration covering the elements the human author contributed to the AI-generated work, including the text and selection and arrangement of the text created by the author).

¹⁵⁰ COMPENDIUM OF COPYRIGHT PRACTICES, *supra* note 15, at § 313.2; U.S. COPYRIGHT OFF., REPORT TO THE LIBRARIAN OF CONGRESS BY THE REGISTER OF COPYRIGHTS 5 (1966).

registration for works created with generative AI is too rigid and ultimately incorrect.¹⁵¹ In his article, Professor Lee suggests that the Copyright Office misunderstands the dynamism of authorship and therefore reduces it to a regimented task that eschews experimentation and creativity.¹⁵² The Office's guidance states authorship for copyright purposes merely relies on whether the underlying work was created by a human.¹⁵³ Professor Lee finds that the Office's position supersedes the requirement that a human contribute to a work; in fact, he notes that it requires potential authors to overly emphasize the premeditation of a specific result and exercise too much control over generated images in order to avoid randomness in their generation.¹⁵⁴ The flaw in Professor Lee's description of the Office's position is that no series of word prompts, no matter how attuned they are to a given AI system, can adequately "dictate a specific result" in the same way as a person applying pen to paper, paintbrush to canvas, or compressing a camera shutter pointed at a particular scene at a particular moment in time.¹⁵⁵

Generative AI does all the conception and execution of the "work" on its own and produces an image or text that corresponds to the prompts it is provided.¹⁵⁶ In other words, generative AI regurgitates mash-ups of the content it ingests and, by way of an algorithm and prompts, generates a menu of related material.¹⁵⁷ No copyrightable human creativity, absent the works on which the AI is trained, is performed in the production of the images or text the AI generates and no human author is involved in this step.¹⁵⁸ Notably, however, this process essentially replaces

¹⁵¹ Lee, *supra* note 96, at 18.

¹⁵² *Id.* at 4.

¹⁵³ See COMPENDIUM OF COPYRIGHT PRACTICES, *supra* note 15, at § 313.2 (stating that copyright registration will be denied to non-human creators, such as animals, nature, plants, or supernatural beings).

¹⁵⁴ Lee, *supra* note 96, at 19.

¹⁵⁵ See *id.* (discussing the controversial flaws of the Copyright Office's position by requiring users to generate results with text prompts that command a certain result in the final creation); Steve Engelbrecht, *Output from AI LLMs is Non-Deterministic. What That Means and Why You Should Care.*, SITUATION (May 12, 2023), <https://www.situation.com/non-determinism-in-ai-llm-output/> (explaining that due to the non-deterministic nature of generative AI models, including their complex neural networks and immense amount of training data, the AI technology can produce different outputs despite being given identical inputs); Katherine Lee, A. Feder Cooper & James Grimmelman, *Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain*, J. OF THE COPYRIGHT SOC'Y OF THE U.S.A. 14 (forthcoming 2024), <https://james.grimmelman.net/files/articles/talkin-bout-ai-generation.pdf> (highlighting that in AI-generative modeling, due to the randomness in the model, it is not unlikely that running the same model with the same inputs will result in different outputs).

¹⁵⁶ Lee, Cooper & Grimmelman, *supra* note 155, at 3–6 (explaining the generative AI creation process to highlight which stages reveal human choices that have legal consequences for copyright, including a discussion of how AI generates new creative works based using a handful of vague, broad phrases).

¹⁵⁷ *Id.* at 5 (stating that new creative works outputted from the AI model are based on statistical patterns in the training dataset that are combined in new ways).

¹⁵⁸ Richard A. Crudo, Ivy Clarice Estoesta & William H. Milliken, *From Warhol to War on HAL: Copyright Infringement and Fair Use as Applied to Artificial Intelligence After the Supreme Court's Warhol Decision*, STERNE KESSLER (Aug. 25, 2023), <https://www.sternekessler.com/news->

23

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

the “creation” step that copyright law aims to protect and sidesteps the requirements for creativity reiterated throughout the Copyright Office’s guidance for examiners and applicants.¹⁵⁹

B. Human Creativity in Prompting AI and the Bridgeman Analogy

Not many outputs of generative AI warrant copyright protection—not even outputs imbued with human creativity at the input stage.¹⁶⁰ The uncopyrightability of only certain works is not a new concept introduced with generative AI; it is a controversial but existing conversation in copyright law, underpinned with case law on different copyrightable mediums such as photography.¹⁶¹

The 1999 *Bridgeman Art Library v. Corel Corp.* decision illustrates how photographing pre-existing images may be analogous to providing prompts to an AI software.¹⁶² In the case, photographers were hired to take professional photographs of preexisting two-dimensional artworks for reproduction in publications such as art history books.¹⁶³ Taking these photographs required professional skill, judgment, and talent.¹⁶⁴ For example, specific lighting and camera angles were used so as to avoid glare, as were special camera filters to

insights/publications/warhol-war-hal-copyright-infringement-and-fair-use-applied-artificial/ (detailing the image-generation process by Stable Diffusion AI into three steps, each of which rely almost completely on the AI tool itself and the mathematical representations given to the training data to produce the output – the only human involvement is entering a text prompt).

¹⁵⁹ See *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, supra note 112 (providing guidance on copyright and AI while highlighting the importance of human authorship and the relationship of the authorship elements with human creativity in the context of generative AI). But see Lidia Ratoi, *Sofia Crespo on Artificial Intelligence to Weave New Worlds*, CLOT MAG. (Apr. 16, 2020), <https://clotmag.com/interviews/sofia-crespo-reimagining-nature-through-the-existing> (arguing that, although AI is designed to mimic human decision-making, AI functions more as a tool than as a creator in her art). Ratoi, *supra*. Crespo says that, in her view, only human beings consume art because only human beings feel emotional about images or sounds, no matter if they are technologically generated or handmade by others. *Id.*

¹⁶⁰ See CHRISTOPHER T. ZIRPOLI, CONG. RSCH. SERV., LSB10922, GENERATIVE ARTIFICIAL INTELLIGENCE AND COPYRIGHT LAW 1–2 (2023) (explaining that the Copyright Office’s guidance states that works containing or utilizing AI-generated material may be copyrighted under some circumstances, such as those demonstrating sufficient human creativity, but copyrightability depends on the nature of the human involvement in the copyright process).

¹⁶¹ See *Bridgeman Art Libr. v. Corel Corp.*, 36 F. Supp. 2d 191, 196–97 (S.D.N.Y. 1999) [hereinafter *Bridgeman Art Reconsideration*] (denying copyright protection to photographs that included human creativity at the input stage, as the photographs at issue were exact copies of those available in the public domain and therefore lacked sufficient originality).

¹⁶² See *id.* (holding that the photographs at issue had no spark of human originality and thus were not copyrightable).

¹⁶³ *Bridgeman Art Libr. v. Corel Corp.*, 25 F. Supp. 2d 421, 423–24 (S.D.N.Y. 1998).

¹⁶⁴ See *id.* at 426 (quoting the UK Privy Council’s observation that it takes great skill, labor, and judgment to produce a high-quality copy of a photograph); *Bridgeman Art Reconsideration*, 36 F. Supp. 2d at 197 (stating the court’s opinion that it may be assumed the reproductive photographs at issue required both skill and effort).

maintain fidelity to the original colors.¹⁶⁵ Nevertheless, the court held that the works the photographers produced were ultimately uncopyrightable “slavish copies” of the original artworks.¹⁶⁶ The court noted that, although most photographs are original in some sense, it is not always the case that they are original enough to warrant copyrightability.¹⁶⁷ Furthermore, the court specifically reasoned that the photographs in question lacked sufficient originality because the photographs reproduced a two-dimensional existing artwork as precisely and with as much fidelity to the pre-existing work as the camera’s technology permits.¹⁶⁸ Although the photographs were inarguably the results of creative *inputs*, the images—or outputs—were meant to look exactly like the original artworks; as such, the visual imprint of the final product was outside the photographers’ creative control.¹⁶⁹

Similarly, although generative AI users may have some ideas as to the output they want to produce, a series of prompts—no matter how skilled, complex, or detailed—do not comprise the copyrightable ingredients of the work that an AI delivers.¹⁷⁰ The AI tool performs the output generation on its own by mathematically manipulating existing creative works.¹⁷¹ This general principle persists regardless of the adjectival numbers or other parameters qualifying a prompt.¹⁷² Ultimately, the image conjured in the author’s mind’s eye differs from the image conjured in the reader’s mind’s eye, and both of these differ from the image generated by the algorithmic process of a given AI model.¹⁷³

To illustrate, envision the following author’s prompt to an AI software: an autumn sunset in Paris on a sunny evening with only a few wispy clouds on a hilly street in the Marais, where three patrons of a sidewalk café are playing cards and drinking absinthe, one patron is wearing a pink shawl, another is smoking a pipe,

¹⁶⁵ *Bridgeman Art*, 25 F. Supp. 2d at 426–27.

¹⁶⁶ *Bridgeman Art Reconsideration*, 36 F. Supp. 2d at 197.

¹⁶⁷ *Id.*

¹⁶⁸ *Id.*

¹⁶⁹ *See id.* (deciding that, despite the skill involved, the photographer’s contribution in this case was analogous to the work done by a photocopier and that simply changing the medium on which the initial creative work is produced cannot, by itself, constitute originality required for copyright).

¹⁷⁰ *See* Second Request for Reconsideration for Refusal to Register SURYAST, *supra* note 80, at 7 (deciding that, although the human author selected the image input, the style of the image, and the specific variable determining how much of the style to transfer, the AI-generated work is not a product of human authorship because the expressive elements of the work were not provided by the applicant and the human had no control over where the stylistic elements would be placed, whether they would be included in the output, and what colors would be used).

¹⁷¹ *See* Lee, Cooper & Grimmelmann, *supra* note 155, at 9 (explaining that AI machine learning algorithms utilize mathematical tools to model patterns in the data and produce content); Alexander Stahl, *The Math Behind Predictions in AI: Unraveling the Magic*, MEDIUM (May 31, 2023), <https://medium.com/@stahl950/the-math-behind-predictions-in-ai-unraveling-the-magic-44b4fcb8af6> (detailing the ways in which mathematics underpins AI models).

¹⁷² *See* Matthew Sag, *Copyright Safety for Generative AI*, 61 HOUS. L. REV. 295, 307 n. 42 (2023) (clarifying that machine learning predictive models emerge from data without explicit foresight or inclusion of relevant parameters and their weights by researchers).

¹⁷³ *See id.* at 327–29 (highlighting the difficulty of prompting generative AI to create pseudo-expressions in pictorial works using two failed copyright infringement provocation examples where the author attempts to produce a near-exact reproduction of an original photograph using prompts, all producing different results).

25

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

and a third is wearing combat boots.¹⁷⁴ Even in the event that the patrons and their immediate surroundings are more thoroughly described by way of descriptive prompting, language is an inadequate tool, on its own, to generate an image.¹⁷⁵ It is the algorithm (in its manifestation at that point in time) coupled with the data set (again, at that point in time) that perform that function.¹⁷⁶ The French prompter will generate a different image than the English prompter, both of whom will generate a different image than the Urdu prompter.¹⁷⁷ The different images do not result from the authors' own inherent creativity, but rather the from generative AI technology's different interpretations of those words.¹⁷⁸

In the *Bridgeman* case, the creativity, skill, and effort involved in photographing the 2-D images culminated in a predetermined and therefore uncopyrightable outcome.¹⁷⁹ With generative AI, the creativity, skill, and effort involved in writing prompts results in a product that the prompter and would-be author ultimately cannot control or reverse engineer. While the logic in this scenario follows a somewhat inverse scenario as compared to the situation in the *Bridgeman* case regarding a predetermined outcome versus a completely unpredictable outcome, the outcome is the same.¹⁸⁰ Creativity by way of input does not necessarily result in creative or copyrightable output.¹⁸¹

In November of 2023, the Authors Guild submitted a comment to the U.S. Copyright Office on a Notice of Inquiry on Artificial Intelligence and Copyright that further emphasized how creative human input does not necessarily translate into original, human-attributable output.¹⁸² The comment noted that an AI system

¹⁷⁴ See *id.* at 299–301 (discussing the ability of generative AI to transform a text prompt inputted by a human into a seemingly infinite number of digital object outputs). The prompt in this sentence was generated by the author of this Article.

¹⁷⁵ See Lee, Cooper & Grimmelmman, *supra* note 155, at 49–50 (discussing the generation step in the creation of generative AI outputs, stating that although users can enter a prompt with highly stylistic description and sophisticated adjectives, it is the parameters embedded in the model that executes operations to perform content generation).

¹⁷⁶ *Id.*

¹⁷⁷ See Snyder, *supra* note 55 (detailing the shortcomings of generative AI's reliability and consistency when using the same model in different languages, as most models are trained on English data and have difficulties properly translating for purposes of content creation).

¹⁷⁸ See *id.* (explaining that generative AI struggles translating underrepresented languages, resulting in unknown and unexplainable results); Lee, Cooper & Grimmelmman, *supra* note 155, at 50, n. 248 (recalling that there are many reasonable and likely outputs for the same input when using generative AI due to the randomness and decoding strategy for language models).

¹⁷⁹ *Bridgeman Art Libr. v. Corel Corp.*, 36 F. Supp. 2d 191, 196–97 (S.D.N.Y. 1999).

¹⁸⁰ See Snyder, *supra* note 55 (highlighting the unpredictability of generative AI's outputs, despite extensive training data and descriptive input prompts).

¹⁸¹ See Second Request for Reconsideration for Refusal to Register SURYAST, *supra* note 80, at 7–8 (deciding that although the human applicant provided inputs to the generative AI software, the work is not a product of human authorship, and therefore not copyrightable, because the expressive elements were in fact generated by the model's interpretation of the inputs and training data without specific contributions or control from the human prompter).

¹⁸² See Authors Guild, Comment Letter on Notice of Inquiry on Artificial Intelligence and Copyright, at 31–32 (Oct. 30, 2023), <https://www.regulations.gov/comment/COLC-2023-0006>.

can generate an infinite number of outputs from any given prompt, but a user of generative AI cannot conceive of or predict the outcome of any particular prompt, no matter its length or descriptiveness.¹⁸³ The submission conceded, however, that there may be cases where the prompts are so detailed that the author directs the AI to produce many iterations and refinements until it can produce what they conceived.¹⁸⁴ In such cases, it may be possible for a human user to have essentially authored the output.¹⁸⁵ Even so, the scenario is perhaps more of a future hypothetical situation than the outcome of any current technological possibility.¹⁸⁶

C. Acknowledging a Sliding Scale for Originality

A sliding scale frequently exists in copyright protection, even absent generative AI.¹⁸⁷ In academic and jurisprudential discourse, this sliding scale is sometimes referred to as “thickness” or “thinness.”¹⁸⁸ The gradations of protection are particularly useful for a discussion of generative AI because there may be windows for “thin” copyright protection.¹⁸⁹

Importantly, copyright protection can apply to works with different degrees of protectiveness in the contexts of technologies like photography and generative AI.¹⁹⁰ The more originality present in a given work, the stronger the copyright protection will be for that work.¹⁹¹ Perhaps this is where copyright’s ability to

9036 (arguing that the Copyright Office should utilize a holistic approach that requires considering various factors of control, conception, execution, and predictability as a way to analyze questions of copyrightability for AI-generated works, and affirming that in some cases human creativity in inputs doesn’t always translate to human authorship that qualifies for copyright protection).

¹⁸³ *Id.*

¹⁸⁴ *Id.* at 32.

¹⁸⁵ *Id.*

¹⁸⁶ *See id.* (arguing that it is only in rare instances that humans can control the precise and final output of generative AI by refining their prompt, as most humans will find it easier and more efficient to utilize other means to control the execution of a work).

¹⁸⁷ *See* Samson Vermont, *The Sine Qua Non of Copyright Uniqueness, Not Originality*, 20 TEX. INTELL. PROP. L.J. 327, 339–42 (2012) (explaining the two ends, and examples in between, of the sliding scale of copyrightability depending on uniqueness).

¹⁸⁸ *See id.* (utilizing ‘thick’ and ‘thin’ as terms to reflect the scale of copyrightability); Dale P. Olson, *Thin Copyrights*, 95 W. VA. L. REV. 147, 156 (1992) (describing how the more original a copyrighted work is, the stronger (“thicker”) its protection will be; on the other hand, the less original, the weaker (“thinner”) the copyright protection will be, eventually verging on being enforceable only against identical replications).

¹⁸⁹ *See* Olson, *supra* note 188, at 148 (noting how some works of authorship may receive minimal copyright protections if they fall into the ‘thin’ side of the copyrightability scale); Velocity of Content, *Digital Hollywood Focus on AI and Copyright*, COPYRIGHT CLEARANCE CTR. (Nov. 13, 2023), <https://velocityofcontentpodcast.com/transcripts/digital-hollywood-focus-on-ai-and-copyright/> (discussing thick versus thin copyrights in the context of photography and analogizing this framework to the use of generative AI).

¹⁹⁰ *See* Velocity of Content, *supra* note 189 (providing high-level discussion about the application of the sliding scale of copyrightability to AI-generated works); Vermont, *supra* note 187, at 339–42 (providing examples of thick vs. thin copyright protection in the context of video footage and photographs).

¹⁹¹ *Id.*

27

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

distinguish between thick copyrights and thin copyrights will play an important role.¹⁹² As technology becomes increasingly sophisticated, and as generative AI itself becomes less of a “black box” through increased explainability, putative authors may be able to demonstrate that they understood an algorithm and the effect of their prompts enough so that they could foresee what the AI would generate.¹⁹³ If an author could demonstrate this understanding, it is possible the Copyright Office could find a sufficient modicum of creativity and therefore deem the work as copyrightable.¹⁹⁴

In Professor Gervais’s recent submission to the U.S. Copyright Office on its Notice of Inquiry on AI and Copyright, he noted that generative AI prompt engineers should generally not be considered authors.¹⁹⁵ Nonetheless, he acknowledged that there are potential exceptions.¹⁹⁶ For example, a series of consecutive detailed prompts containing expressions of specific ideas that reflect human creative choices and are directly perceptible in a machine’s output.¹⁹⁷ In such a case, the prompts’ originality may arguably have been transferred to the output.¹⁹⁸ Professor Gervais still considers such cases the exception to the general rule.¹⁹⁹

Similarly, some argue that providing prompts to an AI machine does not amount to the requisite human intervention sufficient to play a causal role in creative output.²⁰⁰ Advocates of this argument often claim that too much of the

¹⁹² See *Federal Court Upholds Copyright Office Refusal to Register AI Output*, MOSES SINGER (Oct. 11, 2023), <https://www.mosessinger.com/publications/federal-court-upholds-copyright-office-refusal-to-register-ai-output> (discussing the sliding scale of control exercised by a human user of an automated software system in the context of copyrightability).

¹⁹³ See Kernochan Comment, *supra* note 109, at 4 (describing how evaluating whether a human author is an author for copyright purposes is a fact specific inquiry that is poses a difficult analysis given how human action can occur at any and every stage in the AI generation process); Karen McGregor Richmond, Satya M. Muddamsetty, Thomas Gammeltoft-Hansen, Henrik Palmer Olsen & Thomas B. Moeslund, *Explainable AI and Law: An Evidential Survey*, CTR. OF EXCELLENCE FOR GLOB. MOBILITY L. (2023) (offering a suggestion of utilizing explainable AI to better understand AI decision-making and make the technology more predictable and interpretable to humans).

¹⁹⁴ See Kernochan Comment, *supra* note 109, at 4 (highlighting the case-by-case nature of determining human authorship and pointing to the Copyright Office’s requirements that the applicant exercise control over the resulting work to argue there is potential for human authorship over AI-generated works once more is discovered about the functioning of the innerworkings of the AI technologies).

¹⁹⁵ Daniel Gervais, Comment Letter on Notice of Inquiry on Artificial Intelligence and Copyright at 6 (Nov. 1, 2023), <https://www.regulations.gov/comment/COLC-2023-0006-8885>.

¹⁹⁶ *Id.*

¹⁹⁷ *Id.*

¹⁹⁸ *Id.*

¹⁹⁹ See *id.* (explaining the unlikelihood and frequency of AI inputs containing expressions of specific human-created ideas that are specifically and exactly translated to the AI output).

²⁰⁰ See David Newhoff, *Are AI Prompts Authorship in Copyright Law?*, THE ILLUSION OF MORE (Nov. 8, 2022), <https://illusionofmore.com/are-ai-prompts-authorship-in-copyright-law/> (arguing that prompt-writing cannot and should not bestow authorship on the prompter when there is no colorable nexus between the prompt’s text and the selection and arrangement of the work’s creative

generative process is left to chance because a user generally cannot know which sources will inform the machine's response, nor will they be able to control how the algorithm will interact with those sources.²⁰¹ That said, it is possible that with future iterations of the technology, a repeatedly refined set of prompts may sufficiently refine the randomness of the outputs and confer authorship of the result to the user of the technology.²⁰²

Copyright thickness and thinness may become prominent concepts in the AI context as significantly more works are being made with the assistance of generative AI and subsequently registered with the U.S. Copyright Office, which requires copyright applicants to disclaim the elements of their works that were created by (versus with) AI.²⁰³ When an AI technology dictates the expressive elements of its output, the generated material is not the outcome of human authorship and thus must be disclaimed in an application for a copyright registration.²⁰⁴ That said, elements of human creativity may still subsist in a final work.²⁰⁵ The more qualitative and quantitative portions of a work that are created by a human, the thicker the copyright protection will be; the more that is generated by technology, the thinner the copyright protection will be.²⁰⁶ It is perfectly possible and permissible for human (and therefore copyrightable) expression to mingle with technology-generated material, the practical consequence being that the human author must keep track of what they originated.²⁰⁷

elements, as there is no applicable doctrine that would sufficiently harmonize authorship with prompt writing to maintain consistency with the purpose of copyright).

²⁰¹ See Kernochan Comment, *supra* note 109, at 5 (describing the attenuated relationship between a prompter and AI-generated output and arguing that simply providing prompts and making choices offered by the AI technology falls short of sufficient human intervention).

²⁰² See *id.* (acknowledging that technology may evolve in such a way that the question of randomness in AI and attribution to human authorship could be revisited in time).

²⁰³ See *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, *supra* note 112 (indicating that copyright owners must disclaim the portions of the work that were AI-generated in their copyright applications as those components of the ultimate product are not copyright protectable).

²⁰⁴ See *id.* (describing how a machine, not the human is, carrying out the authorship activity when the AI technology receives a prompt from a human and creates a complex work in response); COMPENDIUM OF COPYRIGHT PRACTICES, *supra* note 15, at § 503.5 (stating that copyright registration does not include any unclaimable material and that applicants must exclude that material from the applications).

²⁰⁵ See *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, *supra* note 112 (explaining how inclusion of generative AI assistance does not nullify the possibility that other elements may be copyrightable).

²⁰⁶ *Copyright Protection in AI-Generated Works Update: Decision in Thaler v. Perlmutter*, AUTHORS ALL. (Aug. 24, 2023), <https://www.authorsalliance.org/2023/08/24/copyright-protection-in-ai-generated-works-update-decision-in-thaler-v-perlmutter/>; see Pamela Samuelson, *Why Copyright Law Excludes Systems and Processes from the Scope of Its Protection*, 85 TEX. L. REV. 1921, 1924 (2007) (arguing that it is appropriate to grant thin copyright protection for computer programs; Vermont, *supra* note 187, at 339–42 (explaining that how thickly a work is protected depends on how unique and creative it is)).

²⁰⁷ *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, *supra* note 112. The Copyright Office's Rule provides a helpful example to demonstrate

29

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

III. A MANNION TEST FOR GENERATIVE AI

The outputs of generative AI generally do not possess the requisite level of human creativity to be copyrightable.²⁰⁸ In some cases, outputs of generative AI could benefit from “thin” copyright protection, under which the only recognized form of infringement is virtually identical copying.²⁰⁹ Recent case law provides some guardrails for the aspects of copyrightable creativity in previously disruptive technologies and offers a framework for how courts may assess this sliding scale.²¹⁰ *Mannion v. Coors Brewing Co.*, a 2005 case from the Southern District of New York, explores the contours of creativity necessary to achieve copyrightability in photography.²¹¹ The decision addressed whether a photograph used in a billboard beer advertisement infringed on the copyright of a basketball celebrity’s photograph.²¹² After reviewing the case, the court held that the billboard image’s potential for infringement could not be ruled out, nor ruled in, as a matter of law.²¹³ In reaching its decision, the court also created a three-pronged balancing test to provide guidance on determining whether and to what degree a photograph is copyrightable.²¹⁴

The *Mannion* decision refuted the presumption that because photography generally is copyrightable, all photographs receive copyright protection.²¹⁵ Indeed,

the coexistence of human and technological contribution into one creative work: if a user commands an AI technology to “write a poem about copyright law in the style of William Shakespeare,” they can expect the technology to generate text that is understandable as a poem, discusses copyright, and mimics Shakespeare’s style; but, the AI technology will determine the expressive elements of the output (the rhyming pattern, the text’s structure, and the words in each line). *Id.* Therefore, the generated material is not a result of human authorship, and that material must be disclaimed in the application so as to be excluded from the copyright protection. *Id.*

²⁰⁸ See Second Request for Reconsideration for Refusal to Register *Théâtre d’Opéra Spatial*, *supra* note 7 (rejecting copyrightability claim for generative AI output); Second Request for Reconsideration for Refusal to Register SURYAST, *supra* note 80, at 7 (holding that the AI-generated work is not a product of human authorship because the work’s expressive elements were contributed by the AI technology).

²⁰⁹ See *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, *supra* note 203 (noting how copyrightability may be awarded to works of authorship when disclaiming portions of the work that used AI); *Copyright Protection in AI-Generated Works Update: Decision in Thaler v. Perlmutter*, *supra* note 206 (discussing the implications of the *Thaler* decision regarding the granting of “thin” copyright to AI-generated work altered by a human).

²¹⁰ See *Mannion v. Coors Brewing Co.*, 377 F.Supp.2d 444, 450–54 (S.D.N.Y. 2005) (providing an analysis on human creativity and originality for copyright purposes in the context of photography).

²¹¹ *Id.* (discussing the three respects in which a photograph may be deemed sufficiently original for purposes of copyrightability).

²¹² *Id.* at 447–48.

²¹³ *Id.* at 463.

²¹⁴ *Id.* at 452–54 (establishing a three-part analysis, including rendition, timing, and creation of the subject, for determining the nature and extent of copyright protection that can be granted to a photograph based on its originality).

²¹⁵ See *id.* at 454–55, 463 (highlighting that the parties had catalogued all the similarities and differences between the two photographs at issue in the three-part originality analysis framework and that court found the follow-on work may or may not be an infringement as opposed to a copyrightable work itself).

the court made clear that this presumption is neither true nor practicable.²¹⁶ Nevertheless, the court specified that a human being's creative contributions to a photograph may be copyrightable if the contributions meet the requirements of the *Mannion* test.²¹⁷ The same general recipe can be applied to a copyrightable work for which generative AI was utilized a tool.²¹⁸ Section A of this Part articulates the elements of the Mannion Test, Section B applies these elements to generative AI, and Section C provides additional analysis of originality in generative AI.²¹⁹

A. The Mannion Test

In the *Mannion* decision, Judge Lewis A. Kaplan relied on a British copyright treatise to articulate three general aspects in which a photograph may demonstrate copyrightability: (1) in its rendition; (2) in its timing; and (3) in its creation of the subject.²²⁰ A photograph may be original in any of these three ways, but need not be original in all three cumulatively, to achieve copyrightability.²²¹ It is essential to note that it is not the photographer's interpretation of the work but rather the photograph itself that is the salient point of departure for locating and interpreting the copyrightability factors.²²²

1. Rendition

Judge Kaplan's reference to "rendition" meant that copyright protects not *what* the photograph depicts, but rather *how* the photograph depicts it.²²³ In the

²¹⁶ See *id.* at 450–51, 454, 463 (holding that photographs that fail to satisfy the originality requirements will have limited or nonexistent copyright protection). It is only the human creativity in a work that is eligible for copyright protection. ZIRPOLI, *supra* note 160, at 2. The copyright in a photograph of the Statue of Liberty cannot reach the Statue itself, nor can it cover many elements of the photograph that will be evident in every other photograph of the Statue. See *Mannion*, 377 F.Supp.2d at 450–51 (holding that the photographer of a building or other pre-existing object has no exclusive rights to photographing that subject matter and cannot prevent others from using the same lighting techniques or frames). That said, to the extent the photographer has captured their photograph from a particular angle with particular lighting and other particular effects, their photograph as a whole is copyrightable insofar as that creativity can. See *id.* at 452–55 (holding that copyright protection of a photograph extends only to the original elements in the creation).

²¹⁷ See *id.* at 462–63 (articulating the *Mannion* framework to determine copyrightability in photography).

²¹⁸ See *infra* notes 254–289 and accompanying text.

²¹⁹ See *infra* notes 220–302 and accompanying text.

²²⁰ See *Mannion*, 377 F.Supp.2d at 452–54 (specifying the three-part framework of non-mutually exclusive factors to demonstrate originality for purposes of copyrightability); see also HON. SIR HUGH LADDIE, PETER PRESCOTT & MARY VITORIA, *THE MODERN LAW OF COPYRIGHT AND DESIGNS* (Butterworths, 3d ed. 2000) (providing key tenets of copyrightability).

²²¹ *Mannion*, 377 F.Supp.2d at 452, 454, 463.

²²² See *id.* at 462–63 (focusing the analysis of substantial similarity on the elements in the photographs themselves, not the respective photographers' intentions, talent, skill, or commentary).

²²³ *Id.* at 452 (citing *Caratzas v. Time Life, Inc.*, No. 92 Civ. 6346, 1992 WL 322033 at *4 (S.D.N.Y. Oct. 23, 1992); *Leigh v. Warner Bros.*, 212 F.3d 121, 1214 (11th Cir. 2000)). The *Mannion* decision defines "rendition" to include originality that resides in special features of the photograph, including

Mannion case, the court found that the original photograph and the billboard photograph that resembled it shared a similar composition and angle.²²⁴ Additionally, the two photographs portrayed similar lighting, a cloudy sky backdrop, and subjects wearing similar clothing and jewelry.²²⁵ Yet the court also highlighted differences between the works, including how one image appeared in dim black and white while the other was colorful and bright.²²⁶ Moreover, though similar, the subjects' jewelry was not identical and one subject's T-shirt appeared tighter than the other's.²²⁷ These differences were crucial to the analysis because infringement cannot be found when aspects of dissimilarity exceed the aspects of similarity.²²⁸ Generally, unless one photograph replicates another with total or near-total conformity, the second photograph will be at least somewhat original.²²⁹ Ultimately, the *Mannion* court could not decide whether substantial similarity existed between the two photographs after reviewing the well-catalogued similarities and differences in rendition, and therefore denied summary judgment.²³⁰

In the case of the many photographs of the Statue of Liberty from the southern tip of Manhattan at dusk, the human creativity apparent under the rendition test articulated in the *Mannion* case is protectable.²³¹ The basic features of the image—the Statue, the general lighting qualities, the angle from which the Statue is captured—are only copyrightable to the extent that some creativity is identifiable.²³² Photographers, if interested in claiming copyright in their photographs, may be able to point to a variety of other qualities that differentiate their image from other similar photographs.²³³ Importantly, the creative and original qualities must be discernible in the photographs themselves.²³⁴

Conversely, as discussed in the aforementioned *Bridgeman Art Library* case, no matter how much talent, creativity, or technical skill is invested in the photograph's preparation stage, the photographer's creativity must be reflected in

angle of shot, shade, exposure, light, and effects achieved by use of filters and various techniques.
Id.

²²⁴ *Id.* at 462.

²²⁵ *Id.* at 462–63.

²²⁶ *Id.* at 463.

²²⁷ *Id.*

²²⁸ *Id.* at 463; see also 4 NIMMER ON COPYRIGHT § 13.03[B][1][a], at 13–63.

²²⁹ *Mannion*, 377 F.Supp.2d at 452.

²³⁰ *Id.* at 463.

²³¹ See *id.* (explaining that originality under the “rendition” prong includes the totality of the effect of the precise lighting, angles, and effects given to the photograph through various photographic techniques and filter selection, and thus the copyright analysis must focus on the resulting work's output in final form rather than the inputs).

²³² See *Mannion*, 377 F.Supp.2d at 449, 454 (holding that the extent of copyright protection granted to a photograph will depend on the nature of its originality and restating a long-standing copyright principle that protection can only extend to the original components of a work).

²³³ See *id.* at 462–63 (demonstrating the act of comparing many elements of differentiation between two similar photographs in question to determine copyrightability).

²³⁴ See *id.* (analyzing the photographs holistically and in their final stage, not the photographers' intentions).

the work to qualify for copyright protection.²³⁵ Where “slavish copying” is the ultimate goal, as was in *Bridgeman Art Library*, or where the resulting work is outside putative author’s creative control, as seen in recent applications of generative AI to create works of art, the rendition analysis does not apply because both circumstances result in works that lack sufficiently originality.²³⁶

2. Timing

Judge Kaplan further proposed that, in certain situations (such as Alfred Eisenstaedt’s photograph of a sailor kissing a woman in Times Square on VJ Day or Thomas Mangelsen’s photograph of a brown bear catching a salmon in his mouth) copyright protection may be conferred based on originality in timing.²³⁷ Yet, as with any other aspect of originality, Judge Kaplan underscored that copyright protection based on originality in timing must be limited because copyright in a photograph cannot extend to the subject matter.²³⁸ Indeed, when defining a photographer’s creative expression in the image produced by their camera shutter, the Second Circuit confirmed that timing is a protectable element for photographs insofar as the creative choice exhibited in the moment the photographer captured their subject.²³⁹

Timing is arguably the least creative of the three aspects of originality in a photograph, but it enjoys protection nonetheless.²⁴⁰ Unlike sweat of the brow, the quality of fortuitous—or even meticulously planned—timing cannot be repeated or reproduced, and therefore hews closer to creativity than to skill or effort.²⁴¹

²³⁵ See *Bridgeman Art Library v. Corel Corp.*, 36 F.Supp.2d 191, 196 (S.D.N.Y. 1999) (denying copyright protection to photographic reproductions of two-dimensional public domain artworks despite photographers demonstrating talent and expertise in achieving their ultimate goal of providing faithful reproductions because only a distinguishable variation beyond technical skill will render the reproduction as sufficiently original).

²³⁶ See *id.* at 197 (negating the sweat-of-the-brow doctrine because effort alone does not create the requisite originality or “spark of creativity” required for copyright).

²³⁷ *Mannion*, 377 F.Supp.2d at 452–53.

²³⁸ *Id.* Judge Kaplan stated that the copyright in the photograph of the bear does not protect against later photographs of bears eating salmon in the same location nor does it protect against a similar picture of a different bear catching a different salmon, as copyright doesn’t extend to the natural world the photographer captures. *Id.* at 453.

²³⁹ Jane C. Ginsburg, *U.S. Second Circuit Court of Appeals Tames ‘Transformative’ Fair Use; Rejects ‘Celebrity-Plagiarist Privilege’; Clarifies Protectable Expression in Photographs*, 16 J. OF INTELL. PROP. L. & PRAC. 638, 645 (2021) (commenting on *Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith* (598 U.S. 508 (2023))).

²⁴⁰ See *Mannion*, 377 F.Supp.2d at 453 (describing how timing can warrant copyright protection, but acknowledging the fact that timing creates copyright protection simply because the photographer was “in the right place at the right time”).

²⁴¹ See *id.* at 453 (stating that it requires originality to determine when to photograph something, so as to create the proper effect).

33

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

3. Creation of the Subject

The creation of the subject element focuses on the link between the photographer and the subject matter.²⁴² That is, to the extent the subject of the photograph itself demonstrates creativity, that particular aspect of the photograph may enjoy copyright protection.²⁴³ Judge Kaplan noted that when an artist arranges a scene and then photographs it, they have copyright protection in that scene and can prevent others from duplicating it either via photograph or other mediums.²⁴⁴

This specific category of originality is not typically discussed in the context of copyrightable photography.²⁴⁵ Most photography captures the world as it exists, not as the photographer specifically and purposefully assembles or arranges it for purposes of capturing a particular shot.²⁴⁶ There are, of course, notable exceptions.²⁴⁷ For example, much of the oeuvre of photographer Glen Luchford could be said to fall into this category: he painstakingly sets up the visual compositions that he captures with his lens.²⁴⁸

²⁴² See *id.* at 453–54 (describing how the subject of a photograph may be original to the extent that the photographer created the subject or scene in the photograph).

²⁴³ See *id.* (describing how the court in *Rogers v. Koons* pointed to the original photographer's creativity but neglected to mention that setting up the scene of a couple with puppies in their laps was also original (960 F.2d 301, 304 (2d Cir. 1992))).

²⁴⁴ *Id.*

²⁴⁵ See Amanda Fischer Adian, *Merging Photography's Copyright*, 97 N.Y.U. L. REV. 192, 213 (2022) (noting that most litigation surrounding photography copyrightability relates to ordinary photographs as opposed to photographs created with romantic artistry, such as those at issue in Napoleon Sarony's photography, due to the elements of staging, posing, and arrangement that frequently appear in more commonplace photography).

²⁴⁶ See *Mannion*, 377 F.Supp.2d at 453 (discussing the highly popular photograph of a bear eating salmon in the wild and highlighting that one cannot copyright things occurring in nature); Jonah Berger & Alixandra Barasch, *A Candid Advantage? The Social Benefits of Candid Photos*, 9 Soc. PSYCH. & PERSONALITY SCI. 1010, 1011 (2018) (discussing the growing popularity of candid photography, especially on online social media platforms due to their ability to convey an unvarnished glimpse into how others naturally look and behave); *Candid Photography – Everything You Need to Know*, NFI, <https://www.nfi.edu/candid-photography/> (last visited May 18, 2024) (discussing candid photography's purpose, history, and success in the historical context of photography).

²⁴⁷ See, e.g., Lays Tavares, *Gucci Collectors, a Gucci Campaign That Celebrates "Weird" Obsessions*, L'OFFICIEL (July 7, 2018), <https://www.lofficielmalaysia.com/fashion/gucci-collectors-gucci-fall-winter-2018-campaign-glen-luchford> (describing the purposeful and complex direction and curation of Luchford's photography subjects).

²⁴⁸ See Belle Hutton, *Celebrating 30 Years of Glen Luchford's Fashion Photography*, ANOTHER MAG. (June 26, 2020), <https://www.anothermag.com/fashion-beauty/12626/glen-luchford-online-exhibition-interview-selected-works-1990-2020> (highlighting the unique specificity and detail used by Luchford to establish epic, dramatic, and cinematic photographs through the years).



Photographs by Glen Luchford of Harry Styles in a 2018 Gucci campaign and Gisele Bündchen in a 2023 Louis Vuitton campaign.²⁴⁹

Unlike a spontaneous photograph of the Statue of Liberty from the southern tip of Manhattan at dusk, Mr. Luchford's photographs include a multitude of pre-planned concepts that he captures photographically: these range from the subject's body language or clothing brand, to the unique visual choreography of a given image such as the examples above including piglets or luxurious town cars.²⁵⁰ Arguably, if another photographer took pictures of Harry Styles sporting a tiara and colorful Gucci clothing while carrying a piglet, even if the background and lighting were markedly different, Mr. Luchford may at the very least possess a thin copyright in the "creation of the subject," despite copyright law's well-established idea-expression dichotomy and general denial of protection to ideas.²⁵¹ Mr. Luchford has envisioned and carried out a highly creative visual allegory and that alone is granted some protection, although the individual photograph itself, with *all* its attributes, enjoys a "thick" copyright protection.²⁵² This highlights how

²⁴⁹ Glen Luchford, Photograph of Harry Styles, *in* 2018 GUCCI CAMPAIGN; Glen Luchford, Photograph of Gisele Bündchen, *in* 2018 GUCCI CAMPAIGN. Photographs reproduced with permission by the artist.

²⁵⁰ See, e.g., Michael Beckert, *Glen Luchford Relives 30 Years of Fashion Photography in New Exhibition*, W MAGAZINE (May 29, 2020), <https://www.wmagazine.com/culture/glen-luchford-retrospective-digital-exhibition-interview> (detailing Luchford's extensive and eccentric thought process and instincts when photographing a fashion campaign).

²⁵¹ See 17 U.S.C. § 102 (specifying that ideas are excluded from copyright protection whereas expressions of such ideas warrant protection, otherwise referred to as the Idea-Expression Dichotomy); *Mannion*, 377 F.Supp.2d at 455–56 (holding that the idea-expression dichotomy breaks down in photograph and the elements of a photograph could easily be labeled an "idea" and an "expression," thus the principle does not prevent copyrightability of the subject of photographs).

²⁵² See Robert A. Gorman, *Copyright Protection for the Collection and Representation of Facts*, 76 HARV. L. REV. 1569, 1571 (1963) (describing how certain copyrightable works, such as maps, are virtually absent from creativity but still merit thin copyright protection). Thick copyright protection therefore applies to works that demonstrate creativity in more than one category of creative effort. See Olson, *supra* note 188, at 156 (describing how the thickness of copyright protection increases the more original a copyrighted work is). Mr. Luchford's photographs, for example, exhibit creativity at the stage of creating the subject (staging the photograph) and rendition (capturing the actual photograph and its lighting, angle, etc.). Beckert, *supra* note 250.

35

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

creativity in creation of the subject, while not always applicable, can be a powerful facet of a photograph's overall demonstration of originality.²⁵³

B. Applying the Mannion Test to Generative AI

While the *Mannion* court limited its analysis to the medium of photography, Judge Kaplan identified three criteria for analyzing the creativity in an image: rendition, timing, and creation of the subject.²⁵⁴ Each of these “originality factors” also apply to potential human creativity when using generative AI as a tool.²⁵⁵ Creativity in rendition will be copyrightable when human creativity affects *how* a work looks or *how* a work is written.²⁵⁶ Creativity in timing will be copyrightable when human creativity affects how intimate knowledge of the AI technology used at a certain point in time reflects authorial intention in some way.²⁵⁷ And creativity in creation of the subject will be copyrightable when human creativity achieves an overall output, as generative AI can provide an interesting menu of possibilities for an author to consider, but it is the author who must choose which possibilities to pursue on their own.²⁵⁸ In other words, for each or any of these criteria to apply to a potential claim for copyrightability of a work created using generative AI, human creativity must be both apparent in the specific output for which copyright protection is sought and explainable.²⁵⁹

1. Rendition

The rendition aspect of originality shifts the focus of the copyrightability analysis from the generative AI's output to the human creative expression within it.²⁶⁰ As mentioned in the previous example of café patrons in the Marais in Paris, a generative AI tool serves up a set of images from which a putative author of a single work selects one.²⁶¹ The set of prompts and subsequent selection of an AI-

²⁵³ See *Mannion*, 377 F.Supp.2d at 453–54 (holding that creation of a subject or scene can successfully demonstrate sufficient originality for copyright protection).

²⁵⁴ *Id.* at 450–63 (beginning the analysis with a broad discussion of copyright infringement, but quickly narrowing the discussion to focus on photography as was at issue in the case).

²⁵⁵ See *id.* at 461 (recognizing that the principles in this case may pose a problem in other creative contexts noting, however, that drawing hard lines across areas of the law is difficult and that the categories articulated in the case can be applied elsewhere so long as they are useful and relevant).

²⁵⁶ See *infra* notes 260–272 and accompanying text.

²⁵⁷ See *infra* notes 273–278 and accompanying text.

²⁵⁸ See *infra* notes 279–289 and accompanying text.

²⁵⁹ See *infra* notes 260–289 and accompanying text.

²⁶⁰ See *Mannion*, 377 F.Supp.2d at 452–53 (framing the rendition analysis in photography to focus on the creative expression, such as the lighting, angle, shade, and overall effect of the photograph or output).

²⁶¹ See *Sag*, *supra* note 172, at 324–25 (displaying the fortuitousness of generative AI in providing many similar outputs in response to one input through the example of dystopian photos of an overgrown Chicago cityscape).

generated image has not conferred copyright on this work.²⁶² The author of a work who enjoys copyright by way of rendition may potentially contribute human creativity to the output of generative AI, as is done with technology such as Photoshop, by including several intricate prompts that specify size, style, effects, colors, lighting, exposure, and many other features.²⁶³ In the United States, an author who relies on rendition will also need to disclaim during the copyright registration application process that they utilized generative AI to produce the work, regardless of the amount or type of creativity used in the prompt.²⁶⁴ However, the author may contribute creative elements, such as those specified by Judge Kaplan in the *Mannion* decision, to an AI-generated work before and after the generation of the work that affect *how* it is depicted to the world.²⁶⁵ But, the human creativity will be subservient to the AI generation in this case and the initial image would need to be disclaimed.²⁶⁶

With respect to generative AI, the term “rendition” remains perfectly serviceable as it was defined and explored in *Mannion* because it underscores the importance of the human creativity evident in the potentially copyrightable work.²⁶⁷ If a generative AI technology produces an image or text in response to prompts based on existing material, the output likely cannot benefit from originality in rendition.²⁶⁸ This may change over time as AI technology matures, potentially to the extent putative authors can demonstrate they had thorough and intricate knowledge of the algorithm, how it would respond to prompts, and the author’s

²⁶² See, e.g., Second Request for Reconsideration for Refusal to Register SURYAST, *supra* note 80, at 7 (deciding that, although the human author selected the image and entered a prompt specifying a desired style for the photo output, the AI-generated work is not a product of human authorship and thus does not hold copyright protection).

²⁶³ See *Mr Li v Miss Liu*, Beijing Internet Ct., Case No. 11279, Nov. 27, 2023 (China), at 4–6 (detailing the many specifications and parameters the human specified in the AI prompt, including cartoon styling, bright lighting, soft focus, image size, and subject’s pose); Nicole S. Young, *50 Art Style Prompts for Photoshop Generative AI Fill*, NICOLESY (Sept. 25, 2023), <https://nicolesy.com/2023/09/25/50-art-style-prompts-photoshop-generative-ai-fill/> (describing how new iterations of Photoshop function can assist artists with enhancing generative AI prompting).

²⁶⁴ See *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, *supra* note 112 (stating that copyright applicants must disclaim all portions of the work that were generated by AI in the copyright applications).

²⁶⁵ See *Mannion*, 377 F.Supp.2d at 463 (highlighting that factfinders could reasonably decide a variety of scenarios with respect to photography); Young, *supra* note 263 (discussing the various ways a human can use Photoshop and similar editing technologies to make artistic changes and choices to AI-generated both before and after the work is created by the AI technology).

²⁶⁶ See ZIRPOLI, *supra* note 160, at 2 (stating that only the human-provided components of an AI-generated work can be copyright protected); *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, *supra* note 112 (specifying the disclaimer requirement for any use of AI technology in copyright registration submissions).

²⁶⁷ See *Mannion*, 377 F.Supp.2d at 452–53 (demonstrating that it is a human’s creativity by way of new content, interesting choices, and a variety of other factors that culminates in human creativity).

²⁶⁸ See *id.* (implying that articulation of rendition established in this case cannot extend to mere prompting because any number of people could enter the same inane or intricate prompt; originality for purposes of copyright protection requires something more).

37

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

control over the entirety of the data set on which the algorithm was trained.²⁶⁹ These three basic controls are imperative to support any kind of claim as to human creativity for purposes of rendition considerations, at least with respect to current AI technology.²⁷⁰ That said, the true test of “rendition” laid out for photographs in *Mannion* remains true for an image or text produced by generative AI.²⁷¹ *How* an image is depicted or *how* a text is written constitutes the cornerstone of creativity for purposes of copyright protection under the rendition analysis, and that rendition must emanate from a human being.²⁷²

2. Timing

The timing element that could be copyright protected with respect to AI technology corresponds to putative authors knowing the generative AI algorithm’s mechanisms, probable outputs, and training data set so well *at a particular point in time* that the author could argue they are using AI as a tool comparable to the way a photographer uses a camera to capture a moment in time.²⁷³ An author submitting a work for copyright registration under a theory of originality by way of timing must therefore be in sync with the elements of creativity outlined above and be able

²⁶⁹ See Oguz A. Acar, *AI Prompt Engineering Isn’t the Future*, HARV. BUS. REV. (June 6, 2023), <https://hbr.org/2023/06/ai-prompt-engineering-isnt-the-future> (describing how the technology itself will become more fluent in natural language and prompt engineering may be a fleeting skill); Richmond, Muddamsetty, Gammeltoft-Hansen, Olsen & Moeslund, *supra* note 193 (discussing the use of explainable AI to better understand AI decision processes and assist humans in knowing how to prompt the tool for better results).

²⁷⁰ See *Mannion*, 377 F.Supp.2d at 462–63 (describing how the potentially creative elements of an image should not be analyzed in isolation from each other); Irving Wladawky-Berger, *Why Human Input Matters to Generative AI*, MEDIUM (Oct. 4, 2023), <https://medium.com/mit-initiative-on-the-digital-economy/why-human-input-matters-to-generative-ai-0ed1507fceb2> (discussing the difficulty of ensuring that AI systems are trained on human-made content, as generations of AI models progressively produce less diverse output and lose information about less common aspects of the training data).

²⁷¹ See *Mannion*, 377 F.Supp.2d at 452–53 (emphasizing the importance of how a photograph, the output, is created for purposes of creativity analysis in copyright and defining “rendition” as originality achieved through the artist’s deliberate choices in angle, lighting, exposure, filtering style, and the overall effect that is ultimately manifested in the photograph).

²⁷² See *id.* (detailing the rendition analysis to include considerations only focused on the photographer’s own purposeful and creative contributions to the photograph).

²⁷³ See *id.* at 453 (noting that the image captured due to beneficial timing will be protected only to the identical image, not to any of the underlying elements, as the photographer is accomplishing originality not through creativity expression in the subject through deliberate choice but by capturing a photograph at the right time). This type of analysis is potentially analogous to utilizing a particular data set and token arrangement at a particular point in time to generate a specific output. See *Tokenizer*, OPENAI, <https://platform.openai.com/tokenizer> (last visited May 20, 2024) (explaining the role of tokens, or common sequences of characters in a set of text, in the generative AI process of understanding relationships between words to better predict future outputs).

to demonstrate that the timing was understood and purposefully chosen or harnessed in order to create a specific desired output.²⁷⁴

Just as the photograph of the brown bear capturing the salmon in his mouth receives protection for that particular image at that interesting point in time, so too might AI prompters in the future receive copyright protection for the output of a technological ecosystem for which they have intimate knowledge.²⁷⁵ This may be sufficient to confer a thin copyright.²⁷⁶ Not dissimilar to creativity in timing for photographs, creativity in timing for generative AI will be rare and will challenge the contours of authorial creativity versus phenomena that simply exist and may have been captured with effort and skill but without creativity.²⁷⁷ If and when AI technology evolves so as to enable human authors to have such an intimate knowledge of how an AI will interpret a prompt based on its design and training materials, the timing element may be a relevant consideration for purposes of copyrightability.²⁷⁸

3. Creation of the Subject

²⁷⁴ See *Mannion*, 377 F.Supp.2d at 452–53 (referencing the 1916 case *Pagano v. Chas. Beseler Co.*, in which the court found a photograph of the New York City Library to be protectible for its interesting timing of policemen’s body language and cars waiting for traffic lights, as the court held that it requires originality to simply determine when to take the photograph (234 F. 963 (S.D.N.Y. 1916))).

²⁷⁵ See *id.* (reflecting how the capturing of the fish at that particular moment required intimate knowledge of the ecosystem the photographer was working in); Prashant Gohel, Priyanka Singh & Manoranjan Mohanty, *Explainable AI: Current Status and Future Directions*, IEEE ACCESS (July 12, 2021), <https://arxiv.org/pdf/2107.07045> (outlining the importance of explainable AI for the future of critical applications of the technology, as well as the likelihood and possibilities associated with a human’s understanding of explainable AI in the future).

²⁷⁶ See *Mannion*, 377 F.Supp.2d at 452–53 (describing how others may try to emulate someone’s originality in timing because the thin copyright protection does not extend to the subject or rendition); Justin Ross, *Copyright Cases Visual Artists Should Know: Part 1, Copyrightability*, COPYRIGHT ALL. (Nov. 28, 2023) (articulating that “thin” copyright protection can extend to a natural object at a point in time that expresses requisite minimal creativity through the artist’s copyrightable contributions). Similarly, a thin copyright protection for timing in a generative AI context could only extend to the elements of an output that depended on that author’s keen familiarity with the technological elements of the generative AI tool at a particular point in time. See Frank Hutter, Lin Xu, Holger H. Hoos & Kevin Leyton-Brown, *Algorithm Runtime Prediction: Models & Evaluation*, 206 ARTIFICIAL INTEL. 79, 79 (2014) (explaining that it is possible to predict the length of time it takes an algorithm to run, which has specific application to model building, algorithm analysis, and algorithm selection).

²⁷⁷ See *Mannion*, 377 F.Supp.2d at 452–53 (underscoring how copyright protection based on originality in timing must be limited to works that are truly originality because of the specific moment the photograph was taken).

²⁷⁸ See, e.g., Margaret Beaman, Joanna Tai, Phillip Dawson, David Boud & Rola Ajjawi, *Developing Evaluative Judgement For a Time of Generative Artificial Intelligence*, ASSESSMENT & EVALUATION IN HIGHER EDUC. 1–2 (Apr. 10, 2024), <https://doi.org/10.1080/02602938.2024.2335321> (describing how the intersection between evaluative judgement and generative AI articulates how assessment practices can help users learn to work productively with generative AI).

39

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

Even if photographer Glen Luchford had simply typed “Harry Styles,” “Gucci,” “floral crown,” and “piglets” into a generative AI application, the particular expression of his idea for the subject is singularly expressed in the specific photograph of his subject.²⁷⁹ Generative AI would produce something quite different than Mr. Luchford’s image, although it would still exhibit the prompting elements.²⁸⁰ This is where generative AI tends to fill in blanks that human creators would otherwise need to express on their own.²⁸¹ Whatever prompts are inputted to generate an output, it is the algorithm and the content on which it was trained that ultimately “create” the “subject.”²⁸² The work that an AI technology generates does not benefit from human vision except as can be manifested in a series of words which the human being sprinkles into a machine.²⁸³

The human being then waits for an algorithm to surprise them with its mysterious recipe for creation, which is inexplicable to the human being.²⁸⁴ In the context of existing artistic techniques and technology—be it a paintbrush or a camera—the putative author can modify the image by either repainting or obscuring it, or adjusting the scene or the lighting before capturing it.²⁸⁵ The initial conceptualization of a given work resembles the prompting stage; it is the non-copyrightable preparatory work filled with a fascinating but unrealized stable of ideas.²⁸⁶ Yet assembling the elements of a work in both form and style—even if the artist goes back and edits them before finalizing a work—is the mysterious element

²⁷⁹ See Luchford, *supra* note 249 (demonstrating creativity in the creation of the subject that was reflected in the one specific photograph as manifested by Luchford’s deliberate choices).

²⁸⁰ See Sag, *supra* note 172, at 328–29 (providing two different examples of failed attempts to recreate a specific image using generative AI to show that there is no mathematical guarantee generative AI will produce a close match to the original photograph based on the text prompt).

²⁸¹ See, e.g., Bella Isaacs-Thomas, *How AI Turns Text Into Images*, PBS, <https://www.pbs.org/newshour/science/how-ai-makes-images-based-on-a-few-words> (last updated Jan. 12, 2023) (describing the evolution from GAN networks to current technology that takes natural language and, after incessant training, allows the new models to serve as a background instruction that allows the model to infer concepts like color, objects, and artistic style to ultimately produce an image that matches the textual prompt).

²⁸² *Id.* (specifying that the AI’s algorithm and training data results in the creation of the final product).

²⁸³ See, e.g., Mr Li v Miss Liu, Beijing Internet Ct., Case No. 11279, Nov. 27, 2023 (China), at 4–6 (explaining the various specifications and parameters the human contributed to the prompt, including styling, lighting, focus, size, and pose, which was the only point of human intervention in the creation process of the AI-generated image); *The Power of Prompts in Guiding AI Creativity*, LIGHTON, <https://www.lighton.ai/blog/lighton-s-blog-4/the-power-of-prompts-in-guiding-ai-creativity-36> (last visited May 20, 2024) (providing guidance to users on how to better prompt AI models using creative, well-defined, and specific language to generate more accurate outputs).

²⁸⁴ A. Bandi, P.V.S.R. Adapa & Y.E.V.P.K. Kuchi, *supra* note 53.

²⁸⁵ See, e.g., Aaron A. Agypong, *Photography and Creativity* (2019) (M.S. project, Buffalo State University) (on file with Creative Studies Graduate Student Master’s Projects), <https://digitalcommons.buffalostate.edu/cgi/viewcontent.cgi?article=1307&context=creativeprojects> (emphasizing the essentially human aspect of producing quality photography by thoroughly discussing the various types of photography and the creative elements that characterize them).

²⁸⁶ See *id.* (discussing the way in which photography begins as a mental concept and later evolves into a visual image perceivable by others through utilization of tools like a camera and lighting aids).

of creativity that copyright aims to protect, foster, and disseminate to the world.²⁸⁷ A regurgitation of pre-existing material, even if shuffled in interesting ways through creative prompting, is not copyrightable originality.²⁸⁸ Generative AI will generally fail to meet the “creation of the subject” standard because Glen Luchford’s final expression of his assorted ideas is imbued with the human creativity that any output of generative AI, even when prompted with his unique set of ingredients, will lack.²⁸⁹

C. An Originality Formula for Generative AI

There is no new *Mannion*-like formula necessary for determining human creativity in generative AI outputs because the fundamental requirement of some minimal indicia of human creativity remains intact, even when interpreted for a new technology or artistic genre.²⁹⁰ That said, the *Mannion* factors can be truncated and rephrased more holistically to better frame the copyrightability inquiry in the context of generative AI.²⁹¹ This Article proposes the following standard for AI-assisted works prompted by human beings, based on the principles outlined by Judge Kaplan in the *Mannion* decision:

Does the output of generative AI reflect human originality?²⁹² If so, is the putative author able to explain the process by which they were able to manifest that creativity in the work without relying on the

²⁸⁷ See Mei-lan Stark, *5 Ways Copyright Laws Encourage Personal Expression and Creativity*, U.S. CHAMBER OF COMM. (Apr. 25, 2022), <https://www.uschamber.com/intellectual-property/five-ways-copyright-laws-encourage-personal-expression-and-creativity> (exploring the various ways that copyright law protects and encourages creativity); see, e.g., Roger Beaty, *Why Are Some People More Creative Than Others?*, SCI. AM. (Jan. 16, 2018), <https://www.scientificamerican.com/article/why-are-some-people-more-creative-than-others/> (explaining how neuroscientists are beginning to discern what makes one person more creative than another: creative people are able to co-activate brain networks that normally work separately because the creative brain is ‘wired’ differently); see also Zach Winn, *If Art is How We Express Our Humanity, Where Does AI Fit In?*, MIT NEWS (June 15, 2023), <https://news.mit.edu/2023/generative-ai-art-expression-0615> (referencing *Art and the Science of Generative AI*, and asking whether and how humanity will preserve the human aspect of creativity).

²⁸⁸ See Second Request for Reconsideration for Refusal to Register *Théâtre d’Opéra Spatial*, *supra* note 7, at 3–7 (denying copyright protections to what the U.S. Copyright Office considered to be a recreation of previously created works).

²⁸⁹ See *Mannion*, 377 F.Supp.2d at 453–54 (discussing the limitations to the creation of the subject prong of originality in that a photograph is original only to the extent the human created the scene or subject to be photographed, thus requiring sufficient human creativity in establishing the output).

²⁹⁰ See COMPENDIUM OF COPYRIGHT PRACTICES, *supra* note 15, at § 308.2 (indicating that a slight degree of creativity is required for copyright protections).

²⁹¹ See Grimmelmann, *supra* note 85, at 404 (expressing the author’s opinion that although it is difficult to determine a work’s author when it is created using an unpredictable computer program written by another person, the underlying problems of assigning authorship is no worse in this scenario than elsewhere in copyright).

²⁹² See *Mannion*, 377 F.Supp.2d at 450–54 (presenting the *Mannion* test, which requires human originality in its three focus areas).

41

2024] *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*

skill and labor they expended to do so?²⁹³ In other words, whether by rendition, timing, and/or creation of the subject, can a potential author demonstrate how their personal creativity has manifested itself in a potentially copyrightable work?²⁹⁴

This is perhaps a rearticulation of the tests already proposed by Professors Gervais and Burk, which underscore causation, intent, and volition.²⁹⁵ In other words, this *Mannion*-based “test” is not novel, but simply a more explicit question based in existing copyright jurisprudence.²⁹⁶

As has been true in the United States for almost 150 years, copyright subsists only in photographs when human creativity is apparent in a photograph itself.²⁹⁷ The *Mannion* decision provides some useful contours to help decipher and analyze what, exactly, the elements of originality—and therefore copyrightability—can include.²⁹⁸ The same is true regarding the analysis and consideration of outputs of generative AI, whether they be pictorial, textual, or musical.²⁹⁹ Generative AI is a tool and its outputs at the time of this writing generally reflect too much inexplicable remix and recapitulation of existing copyrighted works to appropriately confer copyright protection to the human being using the tool.³⁰⁰

As a liminal matter, and in the future as AI technology evolves, this general proposition may change.³⁰¹ As generative AI becomes more like a paintbrush or a word processor than a haphazard generator of material that otherwise seems to be copyrightable to human perception, it may become the case that a human prompter is also a presumptive author for purposes of copyright.³⁰²

²⁹³ See *id.* at 451 (noting that effort, skill, and labor do not yield copyright protections, but that protection results from the work’s features themselves).

²⁹⁴ See *id.* at 450–54 (describing the three principles under which to analyze creativity in photography for purposes of copyright protection).

²⁹⁵ See Gervais, *supra* note 100 (presenting a causality copyrightability test for AI-generated outputs); Burk, *supra* note 105 (discussing a copyrightability test for AI-created works rooted in causation, intent, and volition).

²⁹⁶ See *Mannion*, 377 F.Supp.2d at 450–54 (articulating the *Mannion* test to be applied in determining whether a work possesses sufficient originality so as to meet the originality requirements to support copyright protection).

²⁹⁷ See *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 58 (1884) (recognizing the possibility to confer copyright protections to photographs if they are products of human creativity).

²⁹⁸ See *Mannion*, 377 F.Supp.2d at 450–54 (providing an insightful framework to discuss and consider originality of creative works for copyright purposes).

²⁹⁹ See *supra* notes 254–289 and accompanying text.

³⁰⁰ See Ziv Epstein & Aaron Hertzmann, *Art and the Science of Generative AI*, 380 Sci. 1110, 1110–11 (2023) (explaining that generative AI relies on algorithms and numerous opaque variables to yield outputs).

³⁰¹ See Isaacs-Thomas, *supra* note 281 (highlighting how the technology relating to generative AI is rapidly changing over time).

³⁰² See Richmond, Muddamsetty, Gammeltoft-Hansen, Olsen & Moeslund, *supra* note 193 (discussing how explainable AI can eventually allow humans in sufficiently understanding AI decision-making processes, giving humans better control of the technology and creating outputs with more predictability). In the more immediate term, I can imagine that creators for whom

CONCLUSION

Due to the evolving technological landscape and difficulty of truly understanding the innerworkings of artificial intelligence, only human beings should be considered “authors” for purposes of copyright law at this time. Like many new technologies, generative AI has challenged—and will continue to challenge—the continuum that includes the idea-expression dichotomy and the semi-articulate difference between an author and their tool. Not all photographs are copyrightable, although most are to some extent. Similarly, but in reverse, very few generative AI-produced works are copyrightable at this juncture in time. As the technology matures, as people become more adept at manipulating it, and as the content on which generative AI is trained becomes more transparent, there may be stronger and more logical claims that generative AI is a tool and that its users are *bona fide* authors.

For a work that has been made with the assistance of AI, the AI in question must be a tool, not an algorithmic co-author that plumbs an unknowable and potentially infringing ocean of preexisting copyrighted content. Human creativity is not an algorithm trained on existing copyrighted works. It is, I submit, something that the Constitution and Congress understand as important to protect by way of copyright law’s exclusive rights. Until such time as generative AI truly becomes a tool to an author, artificial creativity is an oxymoron.

Recommended Citation: Molly Torsen Stech, *Copyright Thickness, Thinness, and a Mannion Test for Images Produced by Generative Artificial Intelligence Applications*, B.C. INTELL. PROP. & TECH. F. (May 27, 2024), <https://sites.bc.edu/iptf/>.

generative AI also acts as a tool to overcome their physical disabilities may be able to receive authorship status in some circumstances. I include this suggestion as a footnote rather than a discussion point in the body of this paper since it could comprise a new article itself. Generative AI arguably has not reached a stage of maturity to be of great creative assistance to visually impaired people, for example. *See, e.g.*, Laurie Henneborn, *Designing Generative AI to Work for People with Disabilities*, HARV. BUS. REV. (Aug. 18, 2023), <https://hbr.org/2023/08/designing-generative-ai-to-work-for-people-with-disabilities>; Stefan Milne, *Can AI Help Boost Accessibility? These Researchers Rested It for Themselves*, UNIV. OF WASHINGTON NEWS (Nov. 2, 2023), <https://www.washington.edu/news/2023/11/02/ai-accessibility-chatgpt-midjourney-ableist/>. *But see, e.g.*, Marrakesh Treaty to Facilitate Access to Published Works for Persons Who Are Blind, Visually Impaired or Otherwise Print Disabled, June 27, 2013, WIPO, https://www.wipo.int/edocs/pubdocs/en/wipo_pub_marrakesh_flyer.pdf (explaining that visually impaired people are a category of individuals for whom copyright law has strong precedent in introducing special sets of rules).

GUIDELINES FOR SUBMISSIONS TO *THE TRADEMARK REPORTER*

Founded in 1911, *The Trademark Reporter* (TMR) delivers engaging and comprehensive peer-reviewed scholarship on trademarks and complementary intellectual property from practitioners, academics, and judges worldwide.

The TMR welcomes submissions of articles, commentaries, and book reviews at wknox@inta.org.

The complete TMR Submission Guidelines are available at <https://www.inta.org/resources/the-trademark-reporter/tmr-submission-guidelines/>.

THE TMR ONLINE

International Trademark Association: Issues of the TMR beginning with Volume 105, Number 1, 2015, are available to the public at <https://www.inta.org/TMR>; INTA members may log in for complimentary access to the complete archive.

HeinOnline: Issues of the TMR beginning with Volume 1, Number 1, 1911, are available through HeinOnline.

LexisNexis: Issues of the TMR beginning with Volume 31, Number 1, 1941, are available through LexisNexis.

Westlaw: Issues of the TMR beginning with Volume 80, Number 1, 1990, are available through Westlaw.

PRINTED COPIES

Printed copies of the TMR can be ordered from William S. Hein & Co., Inc. (order@wshein.com).