

12-19-2020

The FTC and AI Governance: A Regulatory Proposal

Michael Spiro
spiro.ms@gmail.com

Follow this and additional works at: <https://digitalcommons.law.seattleu.edu/sjteil>



Part of the [Administrative Law Commons](#), [Agency Commons](#), [Data Storage Systems Commons](#), [Digital Communications and Networking Commons](#), [Intellectual Property Law Commons](#), [Internet Law Commons](#), [Other Computer Engineering Commons](#), [Privacy Law Commons](#), [Risk Analysis Commons](#), [Robotics Commons](#), and the [Science and Technology Law Commons](#)

Recommended Citation

Spiro, Michael (2020) "The FTC and AI Governance: A Regulatory Proposal," *Seattle Journal of Technology, Environmental & Innovation Law*: Vol. 10 : Iss. 1 , Article 2.

Available at: <https://digitalcommons.law.seattleu.edu/sjteil/vol10/iss1/2>

This Article is brought to you for free and open access by Seattle University School of Law Digital Commons. It has been accepted for inclusion in Seattle Journal of Technology, Environmental & Innovation Law by an authorized editor of Seattle University School of Law Digital Commons.

The FTC and AI Governance: A Regulatory Proposal

*Michael Spiro**

I. INTRODUCTION

It is difficult to underestimate the level of excitement and optimism about the ability of artificial intelligence (AI) to address the most pressing challenges facing society.¹ The potential to solve a wide range of complex, challenging problems – in such diverse areas as poverty reduction, education, healthcare delivery, and transportation – seem boundless.² In healthcare, for example, doctors are using AI to anticipate medical complications, resulting in more successful treatments, faster recoveries, and lower costs.³ In transportation, “smart vehicles” powered by AI could save thousands of lives each year by preventing accidents to which human drivers are prone.⁴

AI is fast being adopted throughout all sectors of the

*Michael Spiro is corporate counsel at Smartsheet Inc., a cloud-based work management and execution platform, where he handles data privacy agreements and other privacy-related matters. Mr. Spiro has a Juris Doctor from the University of Washington School of Law, an L.L.M. in Innovation and Technology Law from Seattle University School of Law, and more than fifteen years of experience as a federal judicial law clerk for the Western District of Washington. He wishes to thank all of the friends, family and colleagues without whose support he would not be where he is today.

¹ *Preparing for the Future of Artificial Intelligence*, EXECUTIVE OFFICE OF THE PRESIDENT: NATIONAL SCIENCE AND TECHNOLOGY COUNCIL, COMMITTEE ON TECHNOLOGY 5 (Oct. 2016), https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf [<https://perma.cc/NSG8-2X43>]; Brad Smith & Harry Shum, *The Future Computed: Artificial Intelligence and its Role in Society*, MICROSOFT 48 (2018), https://news.microsoft.com/uploads/2018/02/The-Future-Computed_2.8.18.pdf [<https://perma.cc/8DWA-HGYP>].

² Smith et al., *supra* note 1.

³ EXECUTIVE OFFICE OF THE PRESIDENT, *supra* note 1.

⁴ Smith et al., *supra* note 1, at 49.

economy, with AI research and development experiencing rapid growth.⁵ AI already can be found in everyday technologies such as GPS navigation systems, cellphones, and smart home appliances.⁶ At the heart of these technologies are powerful algorithmic processes that make determinations or predictions based on subtle patterns of behavior in the data they are fed and on which they are trained.⁷ Advocates for AI believe that AI can and will produce more accurate and objective decisions, provide far more personalized experiences, and promote economic efficiency by matching consumers with the products and services they most want or need.⁸

Yet the nature of AI itself, and how it is currently being designed and implemented, presents deep normative legal and ethical challenges to our existing social, economic, and political relationships and institutions.⁹ A core aspect of AI systems is that the algorithmic decision-making processes they employ are largely automated and frequently hidden from the public. Yet, the decisions they make can be highly consequential for individual autonomy and opportunity. AI is increasingly making consumers' decisions for them. But because consumers cannot check what criteria an AI system uses, the AI is only as good as the data it is trained.

Programmers, regulators, and the public alike often imbue AI and algorithmic decision-making with an aura of objectivity, accuracy, and fairness. Despite these assertions, humans are involved in the design, training, and application of AI systems; accordingly, human values are embedded at each level of AI design, training, and application. AI often makes errors because the underlying software and algorithms (written by humans) are faulty or the training data (input by humans) is biased. Even though AI can improve our decision making capabilities, these kinds of errors run the risk of creating or reinforcing unfairness and other systemic problems.¹⁰ For example, despite claims that risk assessment tools provide a more objective measure of the likelihood of recidivism, researchers have found the algorithms powering those tools are written in a way that inaccurately identifies black defendants as future criminals more often than

⁵ Alex Campolo, Madelyn Sanfilippo, Meredith Whittaker & Kate Crawford, *AI Now 2017 Report*, AI Now 3 (2017), https://ainowinstitute.org/AI_Now_2017_Report.pdf [<https://perma.cc/MSJ7-NZF3>].

⁶ *Id.*

⁷ *Id.*

⁸ *Id.*

⁹ *Id.*

¹⁰ Robert Brauneis & Ellen P. Goodman, *Algorithmic Transparency for the Smart City*, 20 YALE J.L. & TECH. 103, 129 (2018).

white defendants.¹¹

These risks mean that the public has a strong interest in ensuring that AI systems are designed and implemented justly.¹² Despite widespread consensus on the need to ensure fairness in this area,¹³ to date, regulation of AI has been limited, and the regulations that do exist are mostly sectoral and industry specific. Given the rapid development and increasing deployment of AI, a far more comprehensive and coordinated approach toward AI governance that promotes fairness in the public interest without overburdening innovation is needed.

As an emerging technology, AI poses significant regulatory challenges. The disproportionate allocation of technical knowledge and expertise between industry and regulators compounds these obstacles.¹⁴ These challenges make the traditional, top-down command-and-control approach to regulation ill-suited to AI governance. Instead, a co-regulatory model that combines industry self-regulation and stakeholder involvement with government oversight is much more promising, at least in the near term. While some have called for the creation of a new overarching federal agency to regulate in this area, given the difficulty of enacting any major new legislation that impacts powerful interests, Congressional action or any chance of a new oversight body is unlikely.¹⁵

A federal agency already exists, however, that is both familiar with and experienced in regulating new and emerging technologies and can step in now to fill the AI regulatory hole. Even with its limited resources, the Federal Trade Commission (FTC) has proven adept at working with industry in the area of data protection, closing many of the “gaps” left by the sectoral approach to privacy regulation in the United States.¹⁶ The FTC has done so largely through its broad powers granted under Section 5

¹¹ Julie Angwin & Jeff Larson, *Bias in Criminal Risk Scores Is Mathematically Inevitable, Researchers Say*, PROPUBLICA (Dec. 30, 2016), <https://www.propublica.org/article/bias-in-criminal-risk-scores-is-mathematically-inevitable-researchers-say> [https://perma.cc/2752-DQ2X].

¹² Brauneis et al., *supra* note 10, at 123.

¹³ Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. DAVIS L. REV. 399, 415 (2017).

¹⁴ Abbey Stemler, *Regulation 2.0: The Marriage of New Governance and Lex Informatica*, 1 VAND. J. ENT. & TECH. L. 87, 103 (2016) (quoting Cristie L. Ford, *New Governance, Compliance, and Principles-Based Securities Regulation*, 45 AM. BUS. L.J. 1, 28 (2008)).

¹⁵ Cameron F. Kerry, *Will This New Congress Be the One to Pass Data Privacy Legislation?*, BROOKINGS (Jan. 7, 2019), <https://www.brookings.edu/blog/techtank/2019/01/07/will-this-new-congress-be-the-one-to-pass-data-privacy-legislation/> [https://perma.cc/M2Q2-4U3K].

¹⁶ Woodrow Hartzog & Daniel J. Solove, *The Scope and Potential of FTC Data Protection*, 83 GEO. WASH. L. REV. 2230, 2233 (2015).

of the FTC Act which allows it to regulate “unfair and deceptive practices.” AI and its applications fall under the scope of that language and are thus regulatable by the FTC. Further, Congress can and should increase the FTC’s resources and give it greater rule-making authority. This will allow the FTC to more effectively meet the challenges that AI presents and will continue to present in the future.

II. AI AND ITS GROWING INFLUENCE

AI consists not of a single technology, but a set of algorithmic processes and computational methods and tools employed in a range of tasks – such as speech and voice recognition, vision, language translation, reasoning, and robotics – that attempt to emulate human capabilities.¹⁷ These technologies are being incorporated across every industry.¹⁸ AI technologies, furthermore, are developing rapidly;¹⁹ the development of self-driving cars being just one example.²⁰ Engineers have developed AI programs with the ability to master highly complex games in just a matter of hours.²¹ And AI-powered drones are now making their appearance in our skies.²²

A. *The Drive Behind the Growth of AI*

This trend has been fueled by big improvements in the quality of algorithms, especially in the area of machine learning,²³

¹⁷ Urs Gasser & Virgilio A.F. Almeida, *A Layered Model for AI Governance*, HARV. UNIV. 2 (Nov. 20, 2017), <https://dash.harvard.edu/bitstream/handle/1/34390353/w6gov-18-LATEX.pdf?sequence=1> [<https://perma.cc/9J6C-QQLC>]; Smith et al., *supra* note 4, at 32-33.

¹⁸ Michael Guihot, Anne F. Matthew & Nicolas P. Suzor, *Nudging Robots: Innovative Solutions to Regulate Artificial Intelligence*, 20 VAND. J. ENT. & TECH. L. 385, 395 (2017).

¹⁹ *Id.* at 403-04; Gasser et al., *supra* note 17, at 1.

²⁰ Luke Dormehl & Stephen Edelstein, *Sit Back, Relax, and Enjoy a Ride Through the History of Self-driving Cars*, DIGITAL TRENDS (Feb. 3, 2019), <https://www.digitaltrends.com/cars/history-of-self-driving-cars-milestones/> [<https://perma.cc/RQ5R-QXF9>] (noting that “[s]eemingly within just a few years, autonomous cars have gone from science fiction fantasy to road-bound reality”).

²¹ Samuel Gibbs, *AlphaZero AI Beats Champion Chess Program After Teaching Itself in Four Hours*, THE GUARDIAN (Dec. 7, 2017), <https://www.theguardian.com/technology/2017/dec/07/alphazero-google-deepmind-ai-beats-champion-program-teaching-itself-to-play-four-hours> [<https://perma.cc/H57L-5ZCF>].

²² Angelica Stabile, *AI-powered Drone Racing May Be the Next Big Sport*, FOX BUS. (Oct. 18, 2019), <https://finance.yahoo.com/news/ai-powered-drone-racing-may-221355086.html> [<https://perma.cc/5NMA-B3CJ>].

²³ Machine learning refers to a range of AI techniques that detect patterns within data, which are then used to make automated decisions or predictions. Harry

by increases in the speed and power of networks and cloud computing, and the exponential growth in the ability of the tech industry to capture, store and process massive amounts of consumer data.²⁴ These developments have been enhanced by the significant progress that has been made in machine sensing, perception, and object recognition.²⁵

Much of the demand and economic incentive for developing and employing AI is driven by the desire to monetize data.²⁶ Indeed, the commercial applications for AI seem limitless, with the largest tech companies investing heavily in the potential to commoditize the technology.²⁷ These trends are only likely to accelerate with the next generation of AI systems.²⁸

B. Machine Learning, Complexity, and Unpredictability

As consumer interactions with AI systems increase, those systems will be making more and more decisions vis-a-vis the data they process that can limit consumers' own ability to choose or decide, mostly without their knowledge.²⁹ The range of potential applications "is so wide that it is not clear that any element of human life would remain untouched by machine learning."³⁰ Many of these decisions will be made automatically, with little human oversight or control, which is a major feature that distinguishes AI

Surden, *Values Embedded in Legal Artificial Intelligence*, UNIV. OF COLO. L. 3 (Oct. 18, 2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2932333## [<https://perma.cc/A32B-8TAS>].

²⁴ Campolo et al., *supra* note 5, at 3; STAN. UNIV., ARTIFICIAL INTELLIGENCE AND LIFE IN 2030, ONE HUNDRED YEAR STUDY ON ARTIFICIAL INTELLIGENCE (Sept. 2016), https://ai100.stanford.edu/sites/default/files/ai_100_report_0831fnl.pdf [<https://perma.cc/D2LX-A6EV>].

²⁵ STAN. UNIV., ARTIFICIAL INTELLIGENCE AND LIFE IN 2030, ONE HUNDRED YEAR STUDY ON ARTIFICIAL INTELLIGENCE (Sept. 2016), https://ai100.stanford.edu/sites/default/files/ai_100_report_0831fnl.pdf [<https://perma.cc/D2LX-A6EV>].

²⁶ *Id.*

²⁷ ION STOICA ET AL., UNIV. OF CALIF. AT BERKLEY, A BERKELEY VIEW OF SYSTEMS CHALLENGES FOR AI 1 (Oct. 16, 2017), <https://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-159.pdf> [<https://perma.cc/TBM2-XP78>]; Guihot et al., *supra* note 18, at 402.

²⁸ STOICA ET AL., *supra* note 27.

²⁹ *Id.*

³⁰ Chris Reed, Elizabeth Kennedy & Sara Nogueira Silva, *Responsibility, Autonomy and Accountability: Legal Liability For Machine Learning*, QUEEN MARY UNIV. OF LONDON, SCH. OF LAW 30 (Oct. 17, 2016), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2853462 [<https://perma.cc/Q5B7-HXAE>].

from prior technologies.³¹

Another distinguishing feature of AI is its unpredictability. The behavior of machine learning algorithms, for example, is almost by definition unpredictable.³² Instead of being specifically programmed to deal with a particular problem, these algorithms are programmed to learn how to solve problems in general, being taught how to do so first under controlled conditions, and then in the real world.³³ The algorithms learn patterns and relationships in the data they are fed and trained on, and build models of the processes involved based on those patterns and relationships.³⁴

But how machine learning algorithms “learn” and how they “reason from experience to practice” can seem quite foreign to humans.³⁵ Adding to the complexity and opacity of the machine learning and reasoning process is the algorithms’ ability to improve how they perform over time based on what they have learned.³⁶ Further, machine learning algorithms do not learn or reason as humans do, which makes their decisions even more difficult to predict or explain.³⁷

As a result, the growth of machine learning means we will be entrusting significant aspects of our lives to algorithmic processes that we do not understand.³⁸ Indeed, we may fast be approaching a time when extremely sophisticated algorithms will themselves decide what information is relevant to consider in making a determination based only on “the vaguest of goals,” thereby becoming essentially self-reliant and autonomous.³⁹ In many applications, machine learning algorithms are able to match or surpass human-level performance.⁴⁰ Clearly, machine learning will have significant implications, especially in regard to individual autonomy, dignity, and self-determination.

Machine learning algorithms also can be extremely fast and complex, categorizing and processing large amounts of data in fractions of a second.⁴¹ And they are highly dynamic, constantly

³¹ Matthew U. Scherer, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARV. J.L. & TECH. 353, 362-63 (2016).

³² Andrew Tutt, *An FDA for Algorithms*, 69 ADMIN. L. REV. 83, 87 (2017).

³³ *Id.*

³⁴ Reed et al., *supra* note 30, at 4.

³⁵ Tutt, *supra* note 32.

³⁶ Surden, *supra* note 23, at 3.

³⁷ Tutt, *supra* note 32.

³⁸ *Id.* at 87-88.

³⁹ *Id.* at 98-100.

⁴⁰ *Id.* at 100; STOICA ET AL., *supra* note 27, at 1.

⁴¹ Tutt, *supra* note 32, at 94; *Digital Decisions*, CTR. FOR DEMOCRACY & TECH. 3 (2018), <https://cdt.org/wp-content/uploads/2018/09/Digital-Decisions-Library->

adjusting and changing their “rules” to fit patterns in the data.⁴² In this way, machine learning algorithms can end up making decisions based on criteria that in many cases may not have been explicitly chosen or anticipated, even by the humans who program them.⁴³ Further, algorithms frequently interact with other algorithms, making it extremely difficult in many instances to predict the various factors that influence an algorithm’s behavior and decision-making.⁴⁴

This kind of unpredictable, but useful, behavior that arises in machine learning algorithms is known as “emergence.”⁴⁵ Emergence is a central goal of the AI field because it reduces the amount of time needed to train a learning algorithm.⁴⁶ Thus, rather than having to code every behavior of the algorithm, programmers can instead simply set the algorithm’s goals and then train it to accomplish those goals.⁴⁷ And, because machine learning algorithms learn from their previous behavior, they continue to improve over time, even without the aid of their human operators.⁴⁸

Not only can machine learning algorithms handle far, and many, more complex rules than humans, their emergent nature also can lead to solutions that humans would not have thought of.⁴⁹ For example, machine learning algorithms often are able to find subtle relationships in data that humans tend to overlook or fail to recognize entirely.⁵⁰ They can do so because learning algorithms operate based on correlations between data points, allowing them to uncover reliable, though incidental, relationships in the data, even though those correlations may defy causal explanation.⁵¹ For these reasons, we often do not – and in many instances cannot –

Printer-Friendly-as-of-20180927.pdf [https://perma.cc/H747-J8BW]; Joan Donovan, Robyn Caplan, Jeanna Matthews & Lauren Hanson, *Algorithmic Accountability: A Primer*, DATA & SOCIETY 2 (Apr. 18, 2018), https://datasociety.net/wp-content/uploads/2018/04/Data_Society_Algorithmic_Accountability_Primer_FINAL.pdf [https://perma.cc/DY8W-AL8P] (noting that algorithm can encode “for thousands of variables across millions of data points”).

⁴² Brauneis et al., *supra* note 10, at 131; Tarleton Gillespie, *The Relevance of Algorithms*, MEDIA TECHNOLOGIES 12-13 (Nov. 26, 2012), http://6.asset.soup.io/asset/3911/8870_2ed3.pdf [https://perma.cc/Q6C9-MH8F]; CTR. FOR DEMOCRACY & TECH., *SUPRA* note 41, at 7.

⁴³ CTR. FOR DEMOCRACY & TECH., *supra* note 41.

⁴⁴ Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CALIF. L. REV. 513, 534 (2015) [hereinafter *Robotics*].

⁴⁵ *Id.* at 532.

⁴⁶ *Id.* at 538-40.

⁴⁷ *Id.* at 538-39; Andrew D. Selbst and Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 97 FORDHAM LAW REV. 1085, 1098 (2018).

⁴⁸ *Robotics*, *supra* note 44, at 539.

⁴⁹ *Id.*; Selbst et al., *supra* note 47.

⁵⁰ Selbst et al., *supra* note 47.

⁵¹ Brauneis et al., *supra* note 10, at 130.

know exactly how a machine learning algorithm arrives at its decision, or whether that decision is fair or accurate.

C. *AI as Power Broker and Arbiter of Information Flow*

These algorithms, which have become endemic to much of our social and digital media landscape, are at the heart of AI, driving its technologies and making crucial decisions about what information consumers can see or access across that landscape – and the opportunities that access may open or limit for us – can have serious impact on many aspects of our lives.⁵² Thus, *in many ways, they are acting as society's "new power brokers," making lightning-fast consequential decisions without the benefit of human oversight.*⁵³

For example, even more aspects of our lives are being “quantified and ranked” by algorithms, resulting in a “scored society.”⁵⁴ In addition to recidivism risk assessment tools that score defendants on the basis of their likelihood to offend,⁵⁵ lenders and other financial institutions are turning increasingly to consumer credit scoring models powered by Big Data and machine learning algorithms.⁵⁶ Such scoring can disparately impact vulnerable populations and other marginalized sectors of society.⁵⁷

⁵² Woodrow Hartzog, *Unfair and Deceptive Robots*, 74 MD. L. REV. 785, 807 (2015); CTR. FOR DEMOCRACY & TECH., *supra* note 41, at 3; Meg Leta Jones, *The Ironies of Automation Law: Tying Policy Knots with Fair Automation Practices Principles*, 18 VAND. J. ENT. & TECH. L. 77, 85 (2015); Iyad Rahwan, *Society-in-the-Loop: Programming the Algorithmic Social Contract*, THE MEDIA LAB 4 (July 20, 2017), <https://arxiv.org/pdf/1707.07232.pdf> [<https://perma.cc/35FG-FKWS>]; Nicholas Diakopoulos, *Algorithmic Accountability Reporting: On the Investigation of Black Boxes*, COLUMBIA JOURNALISM SCH., TOW CTR. FOR DIGITAL JOURNALISM 2-3 (Dec. 2013), http://towcenter.org/wp-content/uploads/2014/02/78524_Tow-Center-Report-WEB-1.pdf [<https://perma.cc/B2HH-Z8MU>].

⁵³ Hartzog, *supra* note 52; Iyad Rahwan, *Society-in-the-Loop: Programming the Algorithmic Social Contract*, THE MEDIA LAB 4 (July 20, 2017), <https://arxiv.org/pdf/1707.07232.pdf> [<https://perma.cc/35FG-FKWS>]; Nicholas Diakopoulos, *Algorithmic Accountability Reporting: On the Investigation of Black Boxes*, COLUMBIA JOURNALISM SCH., TOW CTR. FOR DIGITAL JOURNALISM 2-3 (Dec. 2013), http://towcenter.org/wp-content/uploads/2014/02/78524_Tow-Center-Report-WEB-1.pdf [<https://perma.cc/B2HH-Z8MU>].

⁵⁴ Hartzog, *supra* note 52.

⁵⁵ John Lightbourne, *Damned Lies & Criminal Sentencing Using Evidenced-Based Tools*, 15 DUKE L. & TECH. REV. 327, 331-32 (2017); Katherine Freeman, *Algorithmic Injustice: How the Wisconsin Supreme Court Failed to Protect Due Process Rights in State v. Loomis*, 18 N.C. J. OF L. & TECH. 75, 81 (2016).

⁵⁶ Mikella Hurley & Julius Adebayo, *Credit Scoring in the Era of Big Data*, 18 YALE J.L. & TECH. 148, 163 (2016).

⁵⁷ Hartzog, *supra* note 52, at 808.

Learning algorithms are also taking on a bigger role in determining who gets access to resources and who is able to take advantage of crucial opportunities.⁵⁸ Credit scoring is one area where the impact of these algorithms is being felt. Credit risk assessment tools can adversely affect access to economic opportunities because they rely on data (often inaccurate) collected from a myriad of sources in a non-transparent manner, making it difficult for consumers to verify or challenge the scores they produce.⁵⁹ And despite claims of greater objectivity in predicting who should get access to financial resources, there is no certainty such tools are actually designed to predict creditworthiness, and, in fact, may increase existing bias in the lending industry by scoring consumers based on their religious, community, and familial affiliations, as well as such factors as race or gender.⁶⁰

Many AI systems are beginning to directly impact the physical environment.⁶¹ These range from modern flight control systems in commercial aircraft, to accident-prevention safety features in today's self-driving cars, to robot sailboats scouring the ocean for data.⁶² Furthermore, some systems provide sensitive services that, when performed by people, would require training and certification, such as robots that perform surgical procedures by themselves and other autonomous AI systems deployed in areas traditionally requiring board or bar passage from human practitioners and subject to fiduciary obligations like providing legal, healthcare or financial advice.⁶³

Algorithmic decision-making processes also constitute “a key logic governing the flows of information on which we depend,” guiding our understanding of the world and the decisions we make.⁶⁴ As such, they can be understood as “encoded geographies” that “configure and circumscribe” humans and how we live our lives.⁶⁵ In many digital domains, for example, algorithms are the arbiters that define what data is relevant, making

⁵⁸ CTR. FOR DEMOCRACY & TECH., *supra* note 41.

⁵⁹ Hurley et al., *supra* note 56, at 166.

⁶⁰ *Id.* at 167.

⁶¹ Calo, *supra* note 13, at 417.

⁶² *Id.*; John Markoff, *No Sailors Needed: Robot Sailboats Scour the Oceans for Data*, N.Y. TIMES (Sep. 4, 2016), <https://www.nytimes.com/2016/09/05/technology/no-sailors-needed-robot-sailboats-scour-the-oceans-for-data.html> [<https://perma.cc/Q9K3-J28R>].

⁶³ Calo, *supra* note 13, at 417-19.

⁶⁴ Gillespie, *supra* note 42; Finale Doshi-Velez and Been Kim, *Towards a Rigorous Science of Interpretable Machine Learning* 1 (Mar. 2, 2017), <https://arxiv.org/pdf/1702.08608.pdf> [<https://perma.cc/K2P3-FFSR>].

⁶⁵ Lucas D. Introna, *Algorithms, Governance, and Governmentality: On Governing Academic Writing*, LANCASTER UNIV. 26 (2016), http://www.research.lancs.ac.uk/portal/files/84096709/STHV_final_author_version_introna.pdf [<https://perma.cc/4AMX-38DC>].

legitimate certain matters that are worthy of attention, while reducing the legitimacy of others.⁶⁶ Where the algorithms that power social networks and search engines provide users with personalized, and therefore “ultimately skewed,” results, those that filter the news “narrow what we know, surrounding us in information that tends to support what we already believe.”⁶⁷ In essence, AI can and does dictate much of the information we can or cannot find online.

In a very real sense then, the algorithmic decision-making processes that comprise AI act as a “form of governance,” managing individual consumer behavior, allocating resources, and opening up or closing off opportunities.⁶⁸ Yet, while these automated systems are becoming increasingly embedded into our lives, ordering how we behave and interact digitally in various ways, they mostly “are shrouded in secrecy,” operating in the background largely beyond human control.⁶⁹

To the extent the decisions that such systems make are more accurate and objective, they contribute to economic growth and social progress and have a positive impact overall,⁷⁰ some level of opacity and absence of control is not necessarily problematic. Indeed, there is well-founded hope that AI can be more objective than human decision-makers, can augment human judgment in useful ways, and can reduce bias, leading to better outcomes overall.⁷¹ In a number of ways, however, the very nature of AI systems and algorithmic decision-making is problematic, calling into question exactly how accurate, objective, and fair it really may be.

⁶⁶ *Id.* at 26; Joan Donovan, Robyn Caplan, Jeanna Matthews & Lauren Hanson, *Algorithmic Accountability: A Primer*, DATA & SOCIETY 3 (Apr. 18, 2018), https://datasociety.net/wp-content/uploads/2018/04/Data_Society_Algorithmic_Accountability_Primer_FINAL.pdf [<https://perma.cc/DY8W-AL8P>] (“[O]ur digital engagements . . . become part of algorithmic, automated, and artificially intelligent sorting mechanisms that can either target or exclude us.”).

⁶⁷ Megan Anderle, *How Facebook and Google’s Algorithms Are Affecting Our Political Viewpoints*, HUFFINGTON POST (Oct. 15, 2015), https://www.huffpost.com/entry/how-facebook-and-googles-_b_8282612 [<https://perma.cc/2LC6-Z6QU>].

⁶⁸ Brauneis et al., *supra* note 10, at 114.

⁶⁹ Introna, *supra* note 65, at 26-27; Hartzog, *supra* note 52, 807; CTR. FOR DEMOCRACY & TECH., *supra* note 41.

⁷⁰ Diakopoulos, *supra* note 53, at 2; EXECUTIVE OFFICE OF THE PRESIDENT, *supra* note 1, at 39.

⁷¹ Kate Crawford & Meredith Whittaker, *Artificial Intelligence is Hard to See*, MEDIUM 4 (Sept. 11, 2016), <https://medium.com/@katecrawford/artificial-intelligence-is-hard-to-see-a71e74f386db> [<https://perma.cc/9VDQ-XYK3>]; Campolo et al., *supra* note 5, at 4.

III. THE CONCERNS AI RAISES

A major concern with AI systems is that they are becoming integrated into our economic, social, and political landscape without any rigorous study or validation of their accuracy or objectivity.⁷² Because machine learning algorithms have become so prevalent in our information ecosystem and are making decisions in relation to the data they are fed and trained on that are of significant consequence to consumers, how those algorithms are designed, operate, and interact with consumers matters.⁷³ Algorithms not only control how we access information, communicate with others, and learn about the world, but in many ways they lead us “to internalize their norms and priorities” as we navigate the digital environment.⁷⁴

What data an algorithm ultimately considers, what is excluded from its consideration, and how the data the algorithm relies on is made “algorithm ready”⁷⁵ all influence the algorithm’s output or decision. These factors reflect the choices the algorithm’s designers make concerning the criteria by which the algorithm determines what is relevant.⁷⁶ Those criteria, though, are often obscured from the public.⁷⁷ Rarely do algorithm designers reveal how the criteria they employ “are measured, how they are weighed against one another, what other criteria have also been incorporated, and when, if ever, these criteria will be overridden.”⁷⁸

A. Lack of Consumer Awareness

Consumers are all too frequently unaware, or not fully aware, of how companies use machine learning and other algorithms to thoroughly track users’ online activity.⁷⁹ Indeed, although they encounter algorithms every day, consumers often do not realize that algorithms are at work, let alone what it is they actually do.⁸⁰ For example, consumers largely fail to recognize the

⁷² Crawford et al., *supra* note 71.

⁷³ Gillespie, *supra* note 42; Selbst et al., *supra* note 47, at 3.

⁷⁴ Gillespie, *supra* note 42, at 20-21.

⁷⁵ *Id.* at 2-5. Before data can be fed to an algorithm, it must be collected, and then readied or “cleaned up” for the algorithm, which means that sometimes data is excluded or “demoted.”

⁷⁶ *Id.* at 2.

⁷⁷ *Id.*

⁷⁸ *Id.* at 10.

⁷⁹ *Id.* at 7.

⁸⁰ Solon Barocas, Sophie Hood & Malte Ziewitz, *Governing Algorithms: A Provocation Piece*, prepared for the “Governing Algorithms” conference, N.Y. UNIV. 6 (May 17, 2013), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2245322 [<https://perma.cc/C5RA-WFDV>].

extent to which the designers of algorithms are conducting research to understand and “operationalize” how people “habitually seek, engage with, and digest information,” or that the results of this research are then used in the design of the underlying algorithms to better shape user behavior.⁸¹ Even if users are aware of such activity, they often have no ability to challenge it.⁸²

Consumers also are often unaware of *the extent* to which algorithms are employed to track their online activity across linked websites, and the incentives those algorithms are designed to create for consumers to remain within that information ecosystem.⁸³ Algorithms allow the gathering of large amounts of each consumer’s data into an ever increasingly comprehensive profile, which is then used to determine access to information, resources, and opportunities.⁸⁴ However, most consumers do not think much, or at all, about the algorithmic criteria that underlies these processes.⁸⁵ Rather, as long as the decision-making systems provide consumers with what they desire, consumers tend to view the decision-making systems as unproblematic: an answer, a solution, or, more often, simply entertainment.⁸⁶

To the extent consumers are aware of the algorithms in their lives, their view is vague, simplistic, and frequently a fundamental misunderstanding of how algorithms operate.⁸⁷ Algorithms tend to be seen as impartial “stabilizers of trust,” offering “practical and symbolic assurances that their evaluations are fair and accurate, and free from subjectivity, error, or attempted influence.”⁸⁸ No algorithmic system is completely free of human influence, so this view depends on the tech industry’s “carefully crafted fiction” of accuracy and objectivity.⁸⁹ Further, given algorithms’ “aura of mechanistic legal neutrality,” consumers are likely to give more deference to their decisions and

⁸¹ Gillespie, *supra* note 42, at 8.

⁸² *Id.* at 7.

⁸³ *Id.*

⁸⁴ *Id.* at 2, 7.

⁸⁵ *Id.* at 12.

⁸⁶ *Id.*

⁸⁷ *Id.* at 19.

⁸⁸ *Id.* at 2, 13; Joan Donovan, Robyn Caplan, Jeanna Matthews & Lauren Hanson, *Algorithmic Accountability: A Primer*, DATA & SOCIETY 2-3, 8 (Apr. 18, 2018), https://datasociety.net/wp-content/uploads/2018/04/Data_Society_Algorithmic_Accountability_Primer_FINAL.pdf [<https://perma.cc/DY8W-AL8P>] (noting that algorithms appear as “unbiased calculations because they take in objective points of reference and provide a standard outcome,” that they “are attractive because they promise neutrality in decision-making,” and that they are seen as “trusted, objective sources of information”).

⁸⁹ Gillespie, *supra* note 42, at 13.

recommendations.⁹⁰

Machine learning algorithms largely operate behind the scenes of AI decision-making systems.⁹¹ The constant dynamic changes that the algorithms undergo tend to be obscured from the public, making it extremely difficult to assess them for bias.⁹² In addition, it is difficult to challenge the reasoning behind a machine's decision or evaluate the objectivity and accuracy of AI decision-making systems because of the inscrutable nature of machine learning and other AI decision-making systems, a natural consequence of their inherent complexity.⁹³ Consumers may be wholly unaware, for example, of just how precisely tailored search engines are to their preferences.⁹⁴ Nor may they understand that search results for the exact same query can vary quite widely depending on who conducts the search precisely because of how specifically tailored those search engines are.⁹⁵

Consumers are not only generally unaware of the sheer volume of unseen decisions being made about them by algorithms, but they also tend to not comprehend the sheer complexity and scale of these systems.⁹⁶ Nor does the public seem to be aware that there is often a level of uncertainty in the operations of even the most sophisticated machine learning algorithms and, therefore, in the decisions those algorithms make.⁹⁷ The level of opacity that surrounds most machine learning systems obscures their inner workings, either due to proprietary protection considerations or to the emergent nature of the systems' algorithms, making them difficult to understand even by those who design them.⁹⁸

⁹⁰ Brauneis et al., *supra* note 10, at 126.

⁹¹ Gillespie, *supra* note 42, at 12.

⁹² *Id.* at 12-13; Crawford et al., *supra* note 71, at 2-3.

⁹³ Selbst et al., *supra* note 47, at 3, 13; CTR. FOR DEMOCRACY & TECH., *supra* note 41; Surden, *supra* note 23.

⁹⁴ Gillespie, *supra* note 42, at 22.

⁹⁵ *Id.*

⁹⁶ Crawford et al., *supra* note 71, at 2-3; Selbst et al., *supra* note 47, at 3; SOLON BAROCAS, KATE CRAWFORD, MEREDITH WHITTAKER, MADELEINE CLARE ELISH, AARON PLASEK & KADIJA FERRYMAN, THE AI NOW REPORT: THE SOCIAL AND ECONOMIC IMPLICATIONS OF ARTIFICIAL INTELLIGENCE TECHNOLOGIES IN THE NEAR-TERM, AI NOW INST. 23 (July 7, 2016), https://ainowinstitute.org/AI_Now_2016_Report.pdf [<https://perma.cc/KW8G-RFUD>] (a summary of the AI Now public symposium, hosted by the White House and New York University's Information Law Institute); *Ethically Aligned Design: A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems*, THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS 90 (Sept. 2016) [hereinafter IEEE], http://standards.ieee.org/develop/indconn/ec/ead_v1.pdf [<https://perma.cc/6XSF-GN9P>].

⁹⁷ IEEE, *supra* note 96.

⁹⁸ Gillespie, *supra* note 42, at 10; John Danaher, *Is Effective Regulation of AI Possible? Eight Potential Regulatory Problems*, PHILOSOPHICAL DISQUISITIONS 7 (July 7, 2015), <http://philosophicaldisquisitions.blogspot.com/2015/07/is->

B. Value Choices and Bias in AI Design

Because an algorithm's decision is imbued with a level of uncertainty, the decision can be wrong. One reason for error is the training data on which the algorithm learns is biased or unrepresentative of the real world.⁹⁹ The data can also exhibit long-standing structural inequalities,¹⁰⁰ and because machine learning algorithms rely on the ability to analyze useful patterns in large volumes of data, such flaws can be particularly problematic.¹⁰¹ Errors also can result from faulty reasoning, criteria, or factor weighing by the algorithm.¹⁰² Yet, given the "probabilistic and emergent" nature of machine learning algorithms and the opacity of AI systems, the public will rarely become aware of them.¹⁰³

Algorithms are value-laden. The designers of algorithms must make a series of decisions about how they will operate and what they can or cannot do.¹⁰⁴ These decisions include what features the algorithm should have, what criteria to emphasize, what data to use, and how to analyze and present that data.¹⁰⁵ Each of these design decisions reflects a specific set of human choices and value-propositions that determine what the algorithm will consider relevant.¹⁰⁶ Given the ubiquity of AI and machine learning systems, such decisions and choices act as "powerful assertions about how things are and are supposed to be."¹⁰⁷

Because AI research and development is largely driven by industry, the design of AI systems and other algorithmic processes is likely to take on the values of industry.¹⁰⁸ The result is that the rights and values of the individual consumer, such as consent, privacy, and transparency, may be ignored, overlooked, or de-

effective-regulation-of-ai-possible.html [https://perma.cc/KB8L-C43E].

⁹⁹ Brauneis et al, *supra* note 10, at 122; Crawford et al., *supra* note 71, at 6.

¹⁰⁰ BAROCAS et al., *supra* note 96.

¹⁰¹ Surden, *supra* note 23; STOICA ET AL., *supra* note 27, at 4.

¹⁰² Brauneis et al, *supra* note 10, at 123.

¹⁰³ Tutt, *supra* note 32, at 90; Crawford, *supra* note 71, at 3.

¹⁰⁴ Surden, *supra* note 23, at 1.

¹⁰⁵ *Id.*; CTR. FOR DEMOCRACY & TECH., *supra* note 41, at 5.

¹⁰⁶ Diakopoulos, *supra* note 70, at 5; CTR. FOR DEMOCRACY & TECH., *supra* note 41, at 5; Joan Donovan, Robyn Caplan, Jeanna Matthews & Lauren Hanson, *Algorithmic Accountability: A Primer*, DATA & SOCIETY 2-3, 8 (April 18, 2018), https://datasociety.net/wp-content/uploads/2018/04/Data_Society_Algorithmic_Accountability_Primer_FINAL.pdf [https://perma.cc/DY8W-AL8P]. Algorithmic decision-making systems also "can take on unintended values that compete with designed values."

¹⁰⁷ Gillespie, *supra* note 42, at 5.

¹⁰⁸ Campolo et al., *supra* note 5, at 31-32.

emphasized.¹⁰⁹ Further, because of this result, the decisions that go into designing AI algorithms are often made with little to no input from those most affected by them.¹¹⁰

C. Fairness, Expectations, and Control of AI

In addition to the values that algorithms embody, AI systems also impact other values society finds important, such as privacy. Machine learning techniques tend to reduce privacy by their ability to detect patterns within large volumes of data and to make easier the aggregation of what were previously disparate data points.¹¹¹ This capacity to recognize subtle data patterns that humans cannot “threatens to eviscerate the already unstable boundary between what is public and what is private.”¹¹² In this way, AI is increasingly able to provide insights into sensitive or intimate details of consumers’ lives, which consumers themselves may never have intended to share.¹¹³

Another potential concern is that the automated decision-making processes at the heart of today’s AI fundamentally differ from those of humans, leading to the generation of solutions that humans would never have intended or expected.¹¹⁴ Humans tend to expect technology to behave according to the same formal and informal norms to which they hold other humans.¹¹⁵ Machine learning algorithms, however, are not bound by the cognitive limitations that the human brain is subject to.¹¹⁶ For example, humans lack the capacity to analyze much of the information at their disposal, particularly when faced with time constraints.¹¹⁷ Not so for machines.

The computational power of machine learning algorithms means that they are far more capable than their human counterparts of identifying possibilities within the data in a given amount of

¹⁰⁹ *Id.*

¹¹⁰ Selbst et al, *supra* note 47, at 3.

¹¹¹ Surden, *supra* note 23, at 4.

¹¹² Calo, *supra* note 13, at 420-21.

¹¹³ *Id.*

¹¹⁴ Scherer, *supra* note 31, at 364.

¹¹⁵ *The National Artificial Intelligence Research and Development Strategic Plan*, NATIONAL SCIENCE AND TECHNOLOGY COUNCIL, NETWORKING AND INFORMATION TECHNOLOGY RESEARCH AND DEVELOPMENT SUBCOMMITTEE 26 (Oct. 2016), https://www.nitrd.gov/PUBS/national_ai_rd_strategic_plan.pdf [<https://perma.cc/9FU7-SYVQ>].

¹¹⁶ For this reason, examples of “emergence”—or “creativity”—that occur in AI systems may in fact be “something of an illusion, a consequence of the computational resources available to these specialized AI programs combined with AI’s freedom from the cognitive biases that affect humans.” Scherer, *supra* note 31, at 364.

¹¹⁷ *Id.* at 364.

time.¹¹⁸ That is, AI systems can come to conclusions that deviate substantially from those humans might typically expect or find acceptable:

Because AI systems are not inherently limited by the preconceived notions, rules of thumb, and conventional wisdom upon which most human decision-makers rely, AI systems have the capacity to come up with solutions that humans may not have considered, or that they considered and rejected in favor of more intuitively appealing options.¹¹⁹

Yet, it is precisely the ability of machine learning algorithms to generate unique and unexpected solutions to problems that makes their use attractive to industry.¹²⁰

People tend to demand that relationships in the decision-making processes they encounter be intuitive, even though the solution the algorithm comes up with may be more “accurate” or “objective.”¹²¹ This demand reflects the human desire to ensure there is some way to assess whether the basis of the process is sound.¹²² Further, as the decision-making processes of machine learning algorithms become more and more opaque, humans will increasingly cease to have ultimate control over how decisions are made, even though such control is a fundamental assumption behind most laws and regulations.¹²³ This lack of control raises the question of whether AI decision-making systems are beginning to regulate humans, rather than the other way around.¹²⁴

Similarly, the fact that much of the validity of machine learning algorithms relies on correlations between data points, rather than actual causation, holds important implications for the acceptability of the algorithms’ conclusions, at least from a human cognitive perspective.¹²⁵ Those correlations can be highly circuitous, and the algorithm’s reasoning process – despite its complexity – tends to involve both simplifications and generalizations.¹²⁶ For example, algorithmic rules are frequently induced from the behavior of entire populations, which is then

¹¹⁸ *Id.*

¹¹⁹ *Id.* at 365.

¹²⁰ *Id.*

¹²¹ Selbst et al., *supra* note 47.

¹²² *Id.* at 15; Surden, *supra* note 23, at 3.

¹²³ IEEE, *supra* note 96, at 92.

¹²⁴ Barocas et al, *supra* note 80, at 9.

¹²⁵ *Id.* at 6-7; Brauneis, *supra* note 10, at 130.

¹²⁶ Barocas et al., *supra* note 80, at 6; Brauneis, *supra* note 10, at 123.

applied to specific individuals.¹²⁷ This application can result in the situation where what can be known about an individual may now be as much dependent on the behavior of *other* people as on what the particular individual says or does.¹²⁸

Thus, while machine learning algorithms may be more accurate, objective, or fair – at least in a computational sense – a question arises as to whether they should be relied on, or relied on without human oversight, for the many consequential decisions impacting consumers. That is, from a policy perspective, we may find that AI’s decision-making processes are not entirely fair.¹²⁹ The need to ensure fairness is especially important given the general consensus that the public has a strong interest in those processes being employed justly.¹³⁰

The need to ensure fairness is all the more important because machine learning algorithms can worsen and perpetuate problems of unfairness through their ability “to scale” and thus to exert comprehensive and uniform influence on decisions, effectively magnifying any error or bias they may embody.¹³¹ For example, search engine algorithms are very selective about the particular data they consider, yet the decisions they make mask a variety of underlying subjective judgments about what data should be included or excluded, how it should be weighed, and how it should be emphasized or deemphasized.¹³² The possibility of error or bias infecting the decision-making process is always present.¹³³

IV. THE CHALLENGE OF REGULATING AI

To date, very few laws or regulations specifically address the unique challenges that AI use poses.¹³⁴ Looking forward, the legal system lacks the resources needed to make effective regulations that will keep up with AI’s rapid pace of research and development.¹³⁵ The tech industry’s inherent complexities and tendency to reward “time-to-market at all cost” reinforce the pacing and other problems with governing AI effectively.¹³⁶

¹²⁷ Barocas et al., *supra* note 80, at 6-7.

¹²⁸ *Id.* at 7.

¹²⁹ Brauneis, *supra* note 10, at 123.

¹³⁰ *Id.*

¹³¹ *Id.* at 129.

¹³² Gillespie, *supra* note 42, at 25; Brauneis, *supra* note 10, at 130.

¹³³ Gillespie, *supra* note 42, at 25.

¹³⁴ Scherer, *supra* note 31, at 356.

¹³⁵ *Id.* at 363.

¹³⁶ Urs Gasser, *AI and the Law: Setting the Stage*, MEDIUM 5 (June 26, 2017), <https://medium.com/berkman-klein-center/ai-and-the-law-setting-the-stage-48516fda1b11> [https://perma.cc/8KUW-PLUH].

A. AI and the Pacing Problem

In addition to the sheer complexity of the technologies themselves, some of the major challenges to effective AI regulation are the scale, heterogeneity, and autonomous nature of many AI systems.¹³⁷ These challenges are further complicated by the uncertainties about how and in what directions AI will develop in the future.¹³⁸ The possibility that many of the risks that AI poses are likely unknown and may even be unknowable further impedes traditional methods of regulation.¹³⁹ Indeed, the uncertainty around AI's potential makes even categorizing the various risks of implementing AI a complicated and arduous task.¹⁴⁰

A major problem that regulators face regarding AI is that technological developments tend to outpace attempts to regulate them.¹⁴¹ Because of the increasing speed of innovation, the technology often disengages or decouples from regulation.¹⁴² This issue is known as the “pacing” problem, where attempts to “future-proof” legislation result in regulatory disconnect, whereby the adopted regulations end up being either too general or too vague to provide meaningful oversight or guidance.¹⁴³

B. Regulators' Disadvantages: Lack of Knowledge and Resources

Of equal significance is the fact that many regulators do not have the resources to adequately address all of the issues that AI technologies present.¹⁴⁴ This information and resource constraint can be particularly problematic in regard to new technologies, such as AI, because there is a steep learning curve and the ability to engage meaningfully with industry experts is necessary to gain the expertise needed to fully understand and act effectively in response to such advancements.¹⁴⁵ Even with sufficient knowledge and expertise, the speed of innovation in the field makes it difficult for regulators to react in a timely and appropriate manner.¹⁴⁶

Regulators also are at a disadvantage given that most of the world's AI is being developed by a handful of large multinational corporations, whose capabilities in the area far outstrip other

¹³⁷ Gasser et al., *supra* note 17, at 4.

¹³⁸ *Id.*; Guihot et al., *supra* note 18, at 390-91.

¹³⁹ Guihot et al., *supra* note 18, at 414.

¹⁴⁰ *Id.* at 397.

¹⁴¹ *Id.* at 421.

¹⁴² *Id.*

¹⁴³ *Id.*

¹⁴⁴ *Id.* at 390.

¹⁴⁵ *Id.* at 431.

¹⁴⁶ *Id.* at 441.

institutions, including the government.¹⁴⁷ On one hand, this imbalance tends to exacerbate the opacity problem, since private firms are more apt to maintain secrecy over their technologies to safeguard their proprietary interests.¹⁴⁸ These companies' heavy investment in AI research and development also leads to "information asymmetries" between those companies and the regulators, the public, and others seeking to understand the technology.¹⁴⁹

Compounding the problem of information asymmetry is even if regulators are able to obtain information, they likely will be unable to fully understand the technology or appreciate its impacts.¹⁵⁰ Indeed, especially when an emerging technology, like AI, is in its beginning stages, only those directly involved in developing the technology may possess the necessary expertise to adequately assess its risks.¹⁵¹ Further, given demand for such expertise, regulators are less likely to be able to compete with industry for top talent.¹⁵² As a result, regulators may be forced to overly rely on industry when attempting to regulate AI.¹⁵³

C. Sectoral Issues, Autonomy, and Unforeseeability

Another complicating factor is that governance of AI is likely to be overseen by more than one regulator, given that while a particular type of AI may be widely adopted, it will be used by different industries in different ways.¹⁵⁴ Moreover, as AI systems become increasingly integrated and embedded into the social and economic environment, the potential for systemic risk becomes amplified, affecting multiple stakeholders, jeopardizing the effectiveness of traditional regulatory models.¹⁵⁵

The nature of AI research and development also makes effective regulation more difficult. There are three general *ex ante* problems with regulating AI research and development:

- *discreetness* (AI projects can be developed with little physical infrastructure, and without the need for large-scale, integrated institutional frameworks);
- *diffuseness* (AI projects can be carried out

¹⁴⁷ *Id.* at 390, 420; Calo, *supra* note 13, at 406, 408.

¹⁴⁸ Guihot et al., *supra* note 18, at 420.

¹⁴⁹ *Id.* at 421, 425-26.

¹⁵⁰ *Id.* at 421-22.

¹⁵¹ Scherer, *supra* note 31, at 384.

¹⁵² *Id.*

¹⁵³ Guihot, *supra* note 18, at 427.

¹⁵⁴ *Id.* at 414-15.

¹⁵⁵ *Id.* at 416-17.

by a variety of diffuse actors in widely dispersed geographic locations); and

- *discreteness* (parties can, without consciously coordinating with each other, make use of “discrete components and technologies ‘the full potential of which will not be apparent until the components come together’”).¹⁵⁶

Traditional regulators are generally ill-equipped to handle these issues.¹⁵⁷

Machine autonomy and algorithmic unpredictability pose other important regulatory concerns. It is inherently difficult trying to control the actions of autonomous systems, particularly when those systems’ decisions or actions are unforeseeable, even to their designers and operators.¹⁵⁸ Indeed, the legal system is likely to struggle to manage these issues in such a way as to ensure aggrieved parties are adequately compensated when AI technologies cause harm.¹⁵⁹ For instance, the legal system may view the behavior of some machine learning processes as so unforeseeable that it would be unfair to extend liability to their designers for harms they cause, leaving those injured thereby with little recourse.¹⁶⁰ Moreover, because the workings of many AI systems are not visible to the public, it can be hard to detect when an AI system’s decision causes harm, or that the system even made the decision, making the concept of redress practically meaningless.¹⁶¹

A related issue concerns AI systems acting in ways that can make them difficult for humans to control.¹⁶² In the most extreme case, an AI system becomes so much smarter and faster than its human counterparts, that it can no longer be controlled by humans at all.¹⁶³ Flawed programming and design also can lead to loss of “local control,” which occurs when those humans who have the legal responsibility for controlling the AI system are no longer able to do so.¹⁶⁴ Loss of control is especially problematic when the emergent nature of the AI system and the interests of its designers

¹⁵⁶ Guihot et al., *supra* note 18, at 390, 420 (quoting Scherer, *supra* note 31, at 359); Scherer, *supra* note 31, at 356-57, 369; Danaher, *supra* note 98.

¹⁵⁷ Scherer, *supra* note 31, at 356.

¹⁵⁸ *Id.* at 357.

¹⁵⁹ *Id.* at 358.

¹⁶⁰ Scherer, *supra* note 31, at 366.

¹⁶¹ CTR. FOR DEMOCRACY & TECH., *supra* note 41, at 4, 8, 12; Danaher, *supra* note 98.

¹⁶² See CTR. FOR DEMOCRACY & TECH., *supra* note 41.

¹⁶³ *Id.*

¹⁶⁴ *Id.*

are no longer in alignment with each other.¹⁶⁵

Given all of these challenges, new, innovative approaches to regulating AI are needed.¹⁶⁶ The old state-centric, command-and-control regulatory model is no longer adequate.¹⁶⁷ As discussed below, while this does not mean that there is no role for government oversight or that regulators do not have an important part to play—indeed they do—rather it is that the growth of machine learning and the emergence of AI calls for a more inclusive, or co-regulatory, approach.

V. THE FTC AND CO-REGULATION: A MODEL FOR GOVERNING AI

Given its sheer scale, complexity, heterogeneity, and autonomous nature, AI requires a systemic change to how it is regulated.¹⁶⁸ As with other transformative technologies, a new approach is needed to preserve—or, if necessary, reexamine and rebalance—those norms and values that are negatively impacted by AI and that society deems important.¹⁶⁹ Nonetheless, because it is difficult to know in advance what an emerging technology is capable of or exactly how it will be used, issues such as the pacing problem — where regulatory efforts ultimately end up being too general or vague to be effective — can arise.¹⁷⁰ Further, for a new technology, some allowance for flexibility in terms of its design and implementation often is required to permit innovation.¹⁷¹

A. New Regulatory Approaches

Several new regulatory approaches have emerged in recent years to address these issues.¹⁷² One approach, co-regulation, combines industry delegation and self-governance with government oversight, making regulation a “shared responsibility” among relevant stakeholders.¹⁷³ It brings legal treatment of the

¹⁶⁵ Danaher, *supra* note 98.

¹⁶⁶ Gasser et al., *supra* note 17, at 4; Guihot et al., *supra* note 18, at 390.

¹⁶⁷ Gasser et al., *supra* note 17, at 4; Guihot et al., *supra* note 18, at 429–30.

¹⁶⁸ *Robotics*, *supra* note 44, at 552–53; Gasser et al., *supra* note 17, at 3.

¹⁶⁹ *Robotics*, *supra* note 44, at 552–53, 558.

¹⁷⁰ Jones, *supra* note 52, at 82.

¹⁷¹ *Id.*

¹⁷² *Id.*

¹⁷³ Gasser et al., *supra* note 17, at 4 Governance under this approach is delegated to the regulated entities based on the recognition that regulators are too far removed from how the emerging technology operates to comprehend and effectively direct its trajectory. Jones, *supra* note 52, at 103. Such delegation seeks to take advantage of the regulated entity’s technical expertise and judgment. Co-regulation brings the public and private sectors together—allowing industry to be integrated into the regulatory decision-making process—and

technology in line with responsible design choices and practices.¹⁷⁴ This approach also tends to be far more responsive—producing better and quicker outcomes—to emerging technology environments, where there are many variables and heterogeneous actors and technological development is dependent on specific contexts and situations; such is the case with AI.¹⁷⁵

Another advantage of co-regulation is that it is often more effective at fostering technological innovation whereas state-centered regulatory efforts risk overburdening innovators or stifling development prematurely.¹⁷⁶ Co-regulation also allows a period of use and development of the technology during which regulators can discover and identify the particular risks or problems that later regulation can more appropriately address.¹⁷⁷ Simultaneously, government oversight is employed to exert the level of control necessary to protect the public from those risks and harms deemed to be unacceptable or undesirable.¹⁷⁸

B. *The Promise of Co-Regulation*

For government oversight of AI under the co-regulatory approach to be effective, a regulator with the ability and willingness to regulate more intrusively when needed is required.¹⁷⁹ There is some debate as to who that regulator should be. In regard to AI, some experts have proposed an overarching agency that has the power to establish a certification process for AI

establishes mechanisms for ongoing communication between industry and the government. Stemler, *supra* note 14, at 103 (quoting Cristie L. Ford, *New Governance, Compliance, and Principles-Based Securities Regulation*, 45 AM. BUS. L.J. 1, 28 (2008)). Regulators set broad regulatory and policy guidelines and goals for industry to act upon. *Id.* Where possible, the government cooperates with industry to determine the means for achieving those goals but refrains from prescribing the precise means for doing so because industry generally has access to superior information and technical expertise. *Id.*

¹⁷⁴ Jones, *supra* note 52, at 102.

¹⁷⁵ *Id.* at 102-03; Reed et al., *supra* note 30, at 31. Governing through co-regulation helps ensure that all competing interests are properly balanced, and that the necessary tradeoffs between the different values the emerging technology embodies are resolved satisfactorily. Rahwan, *supra* note 52, at 3-5. This is an “[i]terative and cooperative” process, which is key to avoiding hastily adopted, and potentially ill-fitting, top-down regulations. Guihot et al., *supra* note 18, at 415.

¹⁷⁶ Guihot et al., *supra* note 18, at 388.

¹⁷⁷ Reed et al., *supra* note 30, at 30.

¹⁷⁸ Guihot et al., *supra* note 18, at 406. While the government acts as more of an “orchestrator rather than a top-down commander” under the co-regulatory approach, it retains a significant role as the ultimate enforcer of regulatory policy. Stemler, *supra* note 14, at 105; Guihot et al., *supra* note 18, at 428.

¹⁷⁹ Guihot et al., *supra* note 18, at 428.

systems.¹⁸⁰ Another proposal contemplates the creation of an agency tasked with supervising the development, deployment, and use of algorithms.¹⁸¹ A unified federal commission also has been proposed for the field of robotics, and such an option could be applicable to AI in general.¹⁸²

One reason for having a single, overarching agency are the efficiencies that it creates in terms of meeting the complex, systemic regulatory challenges that AI poses.¹⁸³ Other agencies that might exert authority over only a small slice of the AI spectrum will lack the necessary expertise, or motivation to regulate its area of responsibility consistently and effectively.¹⁸⁴ A single federal agency, on the other hand, can develop comprehensive, holistic policies rather than piecemeal regulatory efforts.¹⁸⁵ It can quickly respond to new products, practices, and technological changes with more targeted “granular solutions” that better protect consumers.¹⁸⁶ Furthermore, such an agency probably would be more capable of attracting industry talent and thereby build its own expertise.¹⁸⁷

In the United States, regulation of AI has taken the sector-specific, multiple agency oversight approach.¹⁸⁸ In the near term, there seems to be little movement toward taking steps toward regulating AI more comprehensively.¹⁸⁹ And while Congress could move to create a new, overarching federal agency or to enact omnibus legislation dealing explicitly with AI, legislative inactivity over the past 15 years in the area of privacy – which itself has been significantly impacted by the rise of machine learning and other AI technologies – demonstrates that this prospect is fairly grim.¹⁹⁰

C. *The FTC’s Approach to Regulating Technology*

Yet, the United States is not without an agency that has

¹⁸⁰ Scherer, *supra* note 31, at 393.

¹⁸¹ Tutt, *supra* note 32, at 91.

¹⁸² *Robotics*, *supra* note 44, at 556.

¹⁸³ *Id.* at 556-57; Tutt, *supra* note 32, at 114.

¹⁸⁴ Tutt, *supra* note 32, at 115.

¹⁸⁵ *Id.* at 115, 117.

¹⁸⁶ *Id.* at 116-17.

¹⁸⁷ *Id.* at 117; *Robotics*, *supra* note 44, at 557.

¹⁸⁸ *Artificial Intelligence and Life in 2030*, *supra* note 24, at 44. For example, the use of AI in devices delivering medical diagnostics and treatments is subject to regulation by the Food and Drug Administration, the Federal Aviation Administration governs use of drones in regulated airspace, and financial market use of AI technologies is regulated by the Security Exchange Commission. *Id.*

¹⁸⁹ *Id.* at 46.

¹⁹⁰ IRA S. RUBINSTEIN, THE FUTURE OF SELF-REGULATION IS CO-REGULATION, THE CAMBRIDGE HANDBOOK OF CONSUMER PRIVACY 503, 507 (2018).

both the experience and expertise to step in now to fill the regulatory vacuum: the Federal Trade Commission. For example, the FTC has demonstrated its expertise in the field of privacy, which is an area with a similar history of inconsistent or lack of comprehensive regulation.¹⁹¹ In this field, the FTC has moved to assert its authority when industries have not been subject to regulation due to gaps in the country's sectoral laws.¹⁹² In doing so, the FTC has shown an ability to bring a "layer of coherence" to the regulatory system that has solidified over the years.¹⁹³ The FTC is more than capable of taking on the role of regulatory gap filler and coherent AI policy developer.

Instead of simply imposing top-down, command-and-control rules, the FTC has favored a self-regulatory approach to protecting consumer privacy, gradually developing that approach into a regulatory system that over time has become more robust.¹⁹⁴ In the privacy arena, the FTC has taken the position that protecting consumers through self-regulation is more flexible and cost-effective than direct regulation, which at the same time allows for the pace of technological innovation to continue.¹⁹⁵ Further, the norms the FTC has enforced over the years have been developed in accordance with industry stakeholders and consumer expectations; this enforcement strategy is in line with the co-regulatory model of governance.¹⁹⁶

The FTC thus has been consistent in deferring to industry standards where appropriate, thereby avoiding the "dramatic regulatory lurches" that can accompany traditional top-down command-and-control regulation.¹⁹⁷ Because such standards "dictate what is feasible in industry," the FTC's deference keeps its regulatory efforts from being disconnected from industry practice, and therefore from being arbitrarily implemented.¹⁹⁸ As a result, the FTC's regulatory practices are both politically palatable and tend not to overly burden industry.¹⁹⁹ AI can benefit from a similarly "light but steady regulatory approach," through which the FTC can create an environment for AI to thrive, while also regulating to protect consumers and preserve public trust.²⁰⁰

¹⁹¹ Hartzog et al., *supra* note 16.

¹⁹² *Id.*

¹⁹³ *Id.* at 2269.

¹⁹⁴ Ira S. Rubinstein, *Privacy and Regulatory Innovation: Moving Beyond Voluntary Codes*, 6 J. LAW & POL'Y FOR THE INFO. SOC'Y 355, 356 (2011); Hartzog et al., *supra* note 16., at 2270.

¹⁹⁵ Rubinstein, *supra* note 194.

¹⁹⁶ Hartzog et al., *supra* note 16, at 2270.

¹⁹⁷ Hartzog, *supra* note 52, at 788.

¹⁹⁸ *Id.* at 825, 829.

¹⁹⁹ *Id.* at 829.

²⁰⁰ *Id.* at 788.

D. Section 5 of the FTC Act

The FTC's primary source of regulatory authority is Section 5 of the Federal Trade Commission Act, which prohibits "unfair or deceptive acts or practices in or affecting commerce."²⁰¹ An unfair or deceptive act or practice is statutorily defined: (1) "a material 'representation, omission or practice that is likely to mislead the consumer acting reasonably in the circumstances, to the consumer's detriment'"; or (2) "a practice that 'causes or is likely to cause substantial injury to consumers which is not reasonably avoidable by consumers themselves and not outweighed by countervailing benefits to consumers or to competition.'"²⁰² Section 5 thus bars both "unfair" and "deceptive" trade practices, as well as "unfair methods of competition."²⁰³

This legislative grant of authority is widely recognized as being extremely broad in scope.²⁰⁴ Generally, the FTC has the power to "prosecute any inquiry necessary to its duties in any part of the United States."²⁰⁵ The FTC is further authorized to "gather and compile information concerning, and to investigate from time to time the organization, business, conduct, practices, and management of any person, partnership, or corporation engaged in or whose business affects commerce."²⁰⁶ In addition, the legislative history of Section 5 evinces a clear Congressional intent that the authority the FTC exercises be "evolutionary and wide-reaching."²⁰⁷

Deceptive trade practices include: false representations, sales of hazardous or systematically defective products or services without adequate disclosures, failure to disclose information, use of bait and switch techniques, failure to perform promised services, and failure to meet warranty obligations.²⁰⁸ In terms of such practices, the FTC pursues companies who "exploit consumer ignorance or create a false sense of trust."²⁰⁹ The FTC also has

²⁰¹ *Id.*; Daniel J. Solove and Woodrow Hartzog, *The FTC and the New Common Law of Privacy*, 114 COLUM. L. REV. 583, 599 (2014) (citing 15 U.S.C. § 45(a)(1)).

²⁰² *Id.* (quoting Letter from James C. Miller III to Hon. John D. Dingell, Chairman, House Comm. on Energy & Commerce (Oct. 14, 1983), app. at 174-76; 15 U.S.C. § 45(n)).

²⁰³ Hartzog, *supra* note 52, at 812 (quoting Federal Trade Commission Act, Pub. L. No. 75-447, § 3, 52 Stat. 111 (1938)).

²⁰⁴ *Id.*

²⁰⁵ Solove et al., *supra* note 201, at 609 (quoting 15 U.S.C. § 43).

²⁰⁶ *Id.* (quoting 15 U.S.C. § 46(a)).

²⁰⁷ Hartzog et al., *supra* note 16, at 2246.

²⁰⁸ Hartzog, *supra* note 52, at 792 (quoting Letter from James C. Miller III to Hon. John D. Dingell, app. at 175 (1984)).

²⁰⁹ Solove et al., *supra* note 201, at 666.

shifted to focusing on broken consumer expectations, which incorporate “the universe of preexisting consumer backgrounds, norms, and dispositions,” in addition to elements of design and functionality factors.²¹⁰ The FTC takes consumers “as it finds them,” with all of their cognitive limitations, and prohibits exploitation of those limitations.²¹¹

In considering whether a harm is outweighed by countervailing benefits to consumers or competition in terms of “unfair” trade practices, the FTC takes into account both the cost to consumers to remedy the harm and the cost to society in general.²¹² In essence, this inquiry attempts to separate those instances where consumers are able to protect themselves from those where they are unable to do so.²¹³ If consumers could reasonably have avoided the harm, the FTC will not find a trade practice to be unfair.²¹⁴ Accordingly, most FTC enforcement actions focus on behavior that unreasonably takes advantage of or exploits vulnerable consumers.²¹⁵ However, the FTC also considers whether a trade practice violates established public policy.²¹⁶

The evolutionary and wide-ranging nature of the unfairness standard is the result Congress’s deliberate intention to frame that standard in general terms.²¹⁷ Congress recognized it was not possible to draft a complete set of unfair trade practices, without creating regulatory loopholes or the list quickly becoming outdated.²¹⁸ The breadth of Section 5 is apparent in that it authorizes the FTC to take action against unfair practices that more specific statutes have not yet contemplated.²¹⁹ Indeed, the FTC has the authority to determine a practice is unfair, even if it is otherwise lawful.²²⁰ The result is that the FTC has significant flexibility in addressing new problems.²²¹

The broad scope of authority Section 5 provides is ideal for responding to the challenges posed by new technologies.²²² Partly as a result of this, the FTC has been able to quickly respond to technological change.²²³ In addition, the FTC has shown it is

²¹⁰ *Id.* at 667.

²¹¹ *Id.*

²¹² *Id.* at 639.

²¹³ Dennis D. Hirsch, *That’s Unfair? Or is it? Big Data, Discrimination and the FTC’s Unfairness Authority*, 103 KY. L.J. 345, 354 (2015).

²¹⁴ *Id.*; Solove et al., *supra* note 201, at 639.

²¹⁵ Solove et al., *supra* note 201, at 639; Hartzog et al., *supra* note 16, at 2250.

²¹⁶ Solove et al., *supra* note 201, at 639-40.

²¹⁷ Hartzog, *supra* note 52, at 812-13.

²¹⁸ *Id.*

²¹⁹ *Id.*

²²⁰ *Id.* at 813.

²²¹ *Id.*

²²² *Id.*

²²³ *Id.* at 825.

capable of fostering emerging technologies while protecting consumers, as it has done in response to the rise of the Internet and the Internet of Things.²²⁴ The FTC has been able to do this at least in part because the unfair and deceptive practices standard is largely technology neutral.²²⁵

*E. The FTC can and Should Exercise its
Section 5 Authority*

Of particular relevance to emerging technologies, and AI specifically, the FTC has shown itself to be capable of regulating the communication, organizational, and design aspects of new technologies.²²⁶ It has acted to protect consumers from privacy and other harms, for example, by notifying commercial firms of their obligation not to act unfairly or deceptively in the design, sale, and use of emerging technologies that interact with consumers.²²⁷ In addition to the broad authority to regulate emerging technologies, the FTC's efforts are further enabled to respond to unfair and deceptive trade practices by the diverse set of tools at its disposal.²²⁸

Although much of the FTC's enforcement activity, vis-à-vis emerging technologies, has been principally in the area of privacy and data protection, there is no reason that the FTC cannot also apply its broad Section 5 authority to machine learning and other automated decision-making processes. During its history, the FTC has repeatedly "recalibrated" how emerging technologies are used to deceive or harm consumers.²²⁹ And given its move to assert its authority in regard to the Internet of Things, the FTC does not need any new grant of authority to confront other new technologies.²³⁰ Rather, it is enough if a new technology is used in commerce to harm or mislead consumers.²³¹

Indeed, the FTC has begun to address the issue of algorithms in the privacy context.²³² Further, the many tools the FTC has – including disclosures and design requirements – can help ameliorate the harms that algorithmic decision-making systems pose.²³³ The FTC also has looked to hold commercial

²²⁴ *Id.* at 788, 814, 825.

²²⁵ Calo, *supra* note 13, at 428.

²²⁶ Hartzog, *supra* note 52, at 788.

²²⁷ *Id.*

²²⁸ *Id.* at 812.

²²⁹ *Id.* at 813–14.

²³⁰ *Id.* at 814.

²³¹ *Id.*

²³² *Id.* at 808–09.

²³³ *Id.* at 809. The FTC, for example, can "refine and articulate technology-specific disclosure rules" where needed. *Id.* at 818. And in the privacy arena, the

entities accountable “for design choices that *indirectly* harm consumers.”²³⁴ Because AI often is employed in the backend of systems with no direct consumer interface, this approach offers a potential solution to harms caused by hidden AI. It could also address harms caused by third parties, since those who facilitate “the wrongful conduct of another” will also trigger FTC action under this theory.²³⁵

For a trade practice to be unfair, the harm must be substantial.²³⁶ The harm can be monetary, but it also may encompass unwarranted health and safety risks.²³⁷ Thus, AI technologies that pose such risks can and should meet the unfairness standard.²³⁸ Many algorithmic decision-making processes, however, will not fall under this category of harm. Further, trivial, speculative, and “other more subjective types of harm” generally do not constitute an unfair practice.²³⁹ Since in many cases it may not be clear the exact extent to which a decision made by an AI system has injured a particular consumer, it may be difficult to establish the requisite level of harm.

On the other hand, notions of what constitute an unfair harm continue to evolve, and there is some indication courts may be open to recognizing more subjective, non-monetary harms under Section 5.²⁴⁰ In addition, the FTC has clarified that a small or incremental injury may constitute sufficient injury if it harms a large number of consumers or if it “raises a significant risk of concrete harm.”²⁴¹ And even where harms might be incremental for only a single individual, if those harms pose a collective problem, the FTC may still be able to act on them.²⁴² Further, the FTC may consider “the cost to society in general” in determining whether there are countervailing benefits to consumers or

FTC has begun to embrace “design-based solutions,” which seek “to create or modify a technology, architecture, or organizational structure or procedure ex ante as an attempt to reduce the likelihood of a harm.” *Id.* at 818–19.

²³⁴ *Id.* at 820 (emphasis added).

²³⁵ *Id.* at 820–21. For example, by permitting use of a third party machine learning or other AI application, an entity could be seen to have furnished others “with the means to engage in the unfair practices,” and therefore could itself be found to have engaged in an unfair trade practice. *Id.* at 820 (quoting DesignWare, LLC, F.T.C. File No. 112 3151, No. C-4390, April 11, 2013).

²³⁶ *Id.* at 814.

²³⁷ *Id.*

²³⁸ *Id.*

²³⁹ *Id.* at 814 (quoting FED. TRADE COMM’N, COMMISSION STATEMENT OF POLICY ON THE SCOPE OF THE CONSUMER UNFAIRNESS JURISDICTION (1980) (hereinafter FTC Statement on Unfairness), *reprinted in* Int’l Harvester Co., 104 F.T.C. 949, 1073 (1984)); Solove et al., *supra* note 201, at 639.

²⁴⁰ Hartzog, *supra* note 52, at 814–15.

²⁴¹ *Id.* at 815; Hartzog et al., *supra* note 16, at 2279, 2283 (quoting FTC Statement on Unfairness, *supra* note 239).

²⁴² Hartzog et al., *supra* note 16, at 2283.

competition.²⁴³

The FTC's authority to promulgate rules defining unfair or deceptive acts or practices is limited, and therefore it must enforce its authority indirectly on a case-by-case basis.²⁴⁴ As such, and because it generally lacks the ability to assess civil penalties, the FTC mostly relies on settlements resulting from its enforcement activities to communicate the rules it wants companies to follow.²⁴⁵ In addition, due to staff and budget constraints, the FTC often must rely on informal complaints and self-reporting of potential violations.²⁴⁶ The FTC's Section 5 authority, furthermore, does not extend to non-profit organizations, common carriers, financial institutions, and certain other entities, nor can it regulate harms committed by consumers in non-commercial contexts.²⁴⁷

Despite these limitations, the FTC has a formidable reputation as an enforcement authority, and commercial entities, and their lawyers, pay close attention to its orders and decisions.²⁴⁸ For example, when the FTC issues a complaint, it is published on the FTC's website, which often generates significant attention in the privacy community.²⁴⁹ One reason for this is the fear firms have of the FTC's auditing process, which not only is "exhaustive and demanding," but can last for as long as 20 years.²⁵⁰ As such,

²⁴³ *Id.* Thus, the FTC has the ability to regulate consumer harms that would not otherwise be adequately addressed by the traditional tort system or through other regulatory efforts. Hartzog, *supra* note 52, at 814.

²⁴⁴ RUBINSTEIN, *supra* note 190; "[T]he FTC has only Magnuson-Moss rulemaking authority, which is so procedurally burdensome that it is largely ineffective." Solove et al., *supra* note 201, at 620 (internal footnote omitted).

²⁴⁵ Solove et al., *supra* note 201, at 605, 620–21. For the most part, the FTC can issue fines only in a contempt action for violating a settlement order. *Id.* at 605. Further, because the fines must reflect the actual amount of loss to consumers, when the FTC does issue fines, those fines often are very small in comparison to the gravity of the violation and overall profits of the violator. *Id.*

²⁴⁶ *Id.* at 609.

²⁴⁷ Hartzog, *supra* note 52, at 815; Hartzog et al., *supra* note 16, at 2289. But because most AI is developed by large commercial firms that fall outside those categories, these limitations are not likely to pose much of an issue.

²⁴⁸ Solove et al., *supra* note 201, at 600, 607, 621.

²⁴⁹ *Id.* at 621.

²⁵⁰ *Id.* at 606. Typically what is required is: "the specific detailing of the agreed-upon safeguards to protect consumer information; an explanation of 'how such safeguards are appropriate to respondent's size and complexity, the nature and scope of respondent's activities, and the sensitivity of the covered device functionality or covered information'; an explanation of 'how the safeguards that have been implemented meet or exceed the protections' agreed upon in the consent order; and a certification of the effectiveness of the company's protections by 'a qualified, objective, independent third-party professional, who uses procedures and standards generally accepted in the profession.'" *Id.* (quoting *In re HTC, Am. Inc.*, FTC File No. 122 3049, No. C-4406, at 5 (F.T.C. July 2, 2013) (consent order)).

the FTC settles most of the enforcement actions it initiates.²⁵¹ Firms are motivated to settle with the FTC because they can avoid having to admit any wrongdoing in exchange for taking remedial measures, and thus they also avoid the costs to their reputation from apologizing.²⁵²

Though done by necessity, the rule-making process the FTC engages in with its consent orders and settlement agreements can be of benefit when regulating emerging technologies.²⁵³ For one, it allows the flexibility needed to adapt to new and rapidly changing situations.²⁵⁴ Further, the FTC can wait and see if an industry consensus develops around a particular standard before codifying that rule through its enforcement actions.²⁵⁵ As with the common law, which has long demonstrated the ability to adjust to technological changes iteratively, the FTC's incremental case-by-case approach can help minimize the risks of producing incorrect or inappropriate regulatory policy outcomes.²⁵⁶

In addition to its use of consent orders and settlement agreements, the FTC has created a type of "soft law" by issuing guidelines, press releases, workshops, and white papers.²⁵⁷ Unlike in enforcement actions, where the FTC looks at a company's conduct and sees how its behavior compares to industry standards, the FTC arrives at the best practices it develops for guidance purposes through a "deep and ongoing engagement with all stakeholders."²⁵⁸ As such, not only is the FTC's authority broad enough to regulate the use of emerging technologies such as AI in commerce, but the FTC's enforcement actions also constitute a body of jurisprudence the FTC can rely on to address the real and potential harms that stem from the deployment of consumer-oriented AI.²⁵⁹

Given its broad grant of authority, the regulatory tools at its disposal, and its experience dealing with emerging technologies, the FTC is currently in the best position to take the lead in regulating AI. The FTC's leadership is sorely needed to fill in the remaining – and quite large – gaps in those few sectoral laws that specifically address AI and algorithmic decision-making.²⁶⁰ Several factors make the FTC the ideal agency for this role. First,

²⁵¹ *Id.* at 610.

²⁵² *Id.* at 610, 613.

²⁵³ Hartzog et al., *supra* note 16, at 2265.

²⁵⁴ *Id.*

²⁵⁵ *Id.*

²⁵⁶ Guihot, *supra* note 18, at 418.

²⁵⁷ Solove et al., *supra* note 201, at 625.

²⁵⁸ *Id.* at 626–27 (quoting Email from David Vladeck, Dir., Bureau of Consumer Prot., to authors (Oct. 3, 2013, 1:12 PM).

²⁵⁹ Hartzog, *supra* note 52, at 815–16.

²⁶⁰ *Id.* at 815.

the FTC can use its broad Section 5 powers to respond rapidly and nimbly to the types of unanticipated regulatory issues AI is likely to create.²⁶¹

Second, the FTC has an established history of approaching emerging technologies with “a light regulatory touch” during their beginning stages, waiting to increase its regulatory efforts only once the technology has become more established.²⁶² This approach provides the innovative space needed for new technologies such as AI to develop to their full potential. Thus, as it has in the past, the FTC would focus on disclosure requirements rather than conduct prohibition, and take a case-by-case approach rather than rely on rulemaking.²⁶³ Also, as it has traditionally done, the FTC can hold public events on consumer-related AI and issue reports and white papers to guide industry.²⁶⁴

In other words, the FTC has long taken a co-regulatory approach to regulation, which it can and should proceed to do with AI. As in other emerging technology areas, this will help industry continue to grow and innovate, while allowing for the calibration among all relevant stakeholders of the “appropriate expectations” concerning the use and deployment of AI decision-making systems.²⁶⁵ At the same time, the FTC should use its regulatory powers to nudge, and when necessary, push companies to refrain from engaging in unfair and deceptive trade practices in the design and deployment of AI systems.²⁶⁶ The FTC should also place the onus on firms that design and implement those systems to ensure misplaced or unrealistic consumer expectations about AI are corrected.²⁶⁷

By nudging (or pushing) firms in this way, the FTC can “gradually impose a set of sticky default practices that companies can only deviate from if they very explicitly notify consumers.”²⁶⁸ In terms of disclosure requirements, as it has done in other contexts, the FTC can develop rules and guidelines for “when and how a company must disclose information to avoid deception and protect a consumer from harm,” which can include requiring firms to adopt the equivalent of a privacy policy.²⁶⁹ Given the black box

²⁶¹ *Id.*

²⁶² *Id.* at 827.

²⁶³ *Id.*

²⁶⁴ *Id.* at 827–28.

²⁶⁵ *Id.* at 828.

²⁶⁶ Solove et al., *supra* note 201, at 673.

²⁶⁷ *Id.*

²⁶⁸ *Id.*

²⁶⁹ Hartzog, *supra* note 52, at 816; IEEE, *supra* note 96, at 93. Such a policy could govern “how the AI should be used, who is qualified to use it, what training is required for operators, and what operators and other people can expect from the AI.” *Id.* This would give consumers a more accurate idea of what to expect from the AI. *Id.*

like nature of most algorithmic decision-making processes, there is much that AI developers might have to disclose to prevent those processes from being deemed unfair or deceptive.²⁷⁰

In addition, given its broad authority under Section 5, the FTC is able to address small, nuanced changes in AI design that could adversely affect consumers, but that other areas of law, such as tort, may not be able to adequately handle.²⁷¹ Again, this is important because AI and algorithmic decision-making can pose profound and systemic risks of harm, even though the actual harm to individual consumers may be small or hard to quantify. And as it has done in the area of privacy, the FTC can become the *de facto* federal agency authority charged with protecting consumers from harms caused by AI systems and other algorithmic decision-making processes.²⁷²

The FTC also can, and should, seek to work with other agencies to address AI-related harms, given that the regulatory efforts of other agencies will still occur and be needed in specific sectors or industries, which would impact and be relevant to the FTC's efforts as well.²⁷³ Agency cooperation is essential to ensuring regulatory consistency, accuracy, and efficiency in the type of complex, varied technological landscape that AI presents.²⁷⁴ This should not be a problem as the FTC's Section 5 authority overlaps regularly with the authority of other agencies, and the FTC itself has a history of cooperating with those agencies.²⁷⁵ Further, the FTC can use its experience working with other agencies to build standards and policy consensus within the regulatory community and among stakeholders.²⁷⁶

The overarching role the FTC has played in protecting consumer privacy within the United States also has given it legitimacy within the wider privacy community. The FTC has been pivotal over time in promoting international confidence in the United States' ability to regulate privacy by for example acting as the essential mechanism for enforcing the Safe Harbor Agreement with the European Union.²⁷⁷ As it takes on a similar overarching regulatory role for AI and algorithmic decision-making processes in this country, the FTC should gain a similar level of legitimacy internationally. This is important given the increasingly cross-

²⁷⁰ Hartzog, *supra* note 52, at 807.

²⁷¹ Hartzog, *supra* note 52, at 820.

²⁷² Solove et al., *supra* note 201, at 600.

²⁷³ Hartzog, *supra* note 52, at 830.

²⁷⁴ *Id.* at 831.

²⁷⁵ *Id.* at 825–26, 830.

²⁷⁶ *Id.* at 825–26.

²⁷⁷ Hartzog et al., *supra* note 16, at 2233, 2270; Solove et al., *supra* note 201, at 604. “Under the Safe Harbor Agreement, companies had to agree to be subject to FTC enforcement authority if they violated [the Agreement’s] principles.” *Id.*

border nature of AI research and development.

So far, the FTC has not pushed the full extent of its authority under Section 5.²⁷⁸ Under a co-regulatory model such as that followed by the FTC, firms are much more willing and motivated to commit to voluntary codes of conduct if faced with a credible threat of stricter government regulation if they fail to abide by the agreed upon standards.²⁷⁹ Thus, the FTC should consider actively pursuing strategies to enforce its unfair and deceptive trade practices authority in the realm of AI to incentivize companies to participate more fully in the co-regulatory process.²⁸⁰

The FTC also has the authority to expand its regulatory reach under Section 5, and should consider new strategies for doing so.²⁸¹ Indeed, the FTC can and should “push in bolder and more aggressive directions,” in light of the risk of harm that AI is likely to continue to pose.²⁸² Because of the nimbleness of the FTC, it is the ideal federal agency to regulate the AI landscape, which, at this point in time, is subject to much uncertainty.²⁸³ Indeed, as one observer has noted, “[n]o other agency has such a broad scope of power over so many different industries,” given that any industry where consumers can be found is subject to the FTC’s enforcement power.²⁸⁴

As one example, the FTC can and should encourage the AI industry to develop codes of conduct for AI research and development and algorithmic decision-making, and then enforce those codes under Section 5’s deception prong, treating a firm’s failure to abide by the relevant code they have agreed to abide by as “a presumptive violation.”²⁸⁵ To increase firms’ incentive to participate in this process, the FTC should also decline to give favorable treatment in enforcement matters to those firms that refuse to abide by the developed codes.²⁸⁶ Further, the FTC should treat all adopted and approved codes of conduct as having established industry standards, thus making them a relevant factor in determining the issue of “reasonableness” under Section 5’s unfairness prong.²⁸⁷

This is not to say that additional legislation by Congress to address the risks and challenges posed by AI decision-making

²⁷⁸ Solove et al., *supra* note 201, at 666; Hartzog et al., *supra* note 16, at 2293-94.

²⁷⁹ RUBINSTEIN, *supra* note 190.

²⁸⁰ *Id.*

²⁸¹ Hartzog et al., *supra* note 16, at 2234.

²⁸² *Id.*

²⁸³ Hartzog, *supra* note 52, at 789.

²⁸⁴ Hartzog et al., *supra* note 16, at 2236.

²⁸⁵ RUBINSTEIN, *supra* note 190.

²⁸⁶ *Id.*

²⁸⁷ *Id.*

systems is not needed. For example, given that academic and other non-profit institutions are involved in AI research and development as well, the FTC should be given authority over non-commercial entities whose practices also harm consumers.²⁸⁸ Further, granting “explicit rulemaking authority” to the FTC would allow it to take a more systematic approach to regulating AI.²⁸⁹ Short of that, though, and rather than rushing to hastily craft new legislation that may be ill-suited to the challenges of regulating AI, Congress should ensure the FTC first has the resources – including the necessary technical expertise – it needs to exert its existing authority effectively.²⁹⁰

VI. CONCLUSION

AI undoubtedly holds enormous potential for addressing many of society’s greatest challenges and will likely bring substantial benefits to consumers and the public at large. For these reasons, the ability to continue to innovate and develop beneficial AI technologies should not be unduly hampered. At the same time, AI, and algorithmic decision-making processes in general, pose significant risks and challenges of their own. Thus, while the most effective method of regulating AI, at least in the near term, is a co-regulatory approach in which the government, industry, and other stakeholders come together to work out appropriate governance standards, there still needs to be an effective oversight body that can step in when and where necessary to enforce those standards.

Although Congress could create a new, overarching federal agency tasked with the regulation of AI, it has shown little inclination over the years to enact comprehensive legislation dealing with this or other similar emerging technology issues. Perhaps, as the industry develops and matures, this will be the path Congress chooses to take. However, in the meantime, the FTC is more than capable of using its broad authority under Section 5 to prevent the tech industry from engaging in unfair and deceptive trade practices in the design and deployment of AI technologies. The FTC can and should use the full extent of that authority to do so. Further, Congress should give the FTC the resources and rule-making authority necessary to effectively regulate in this area. In the meantime, the FTC’s central regulatory role offers the best path forward to govern AI safely and effectively.

²⁸⁸ Hartzog et al., *supra* note 16, at 2286.

²⁸⁹ *Id.* at 2299.

²⁹⁰ *Id.* at 2245 (quoting Szoka Statement on Facebook FTC Privacy Settlement, TECHFREEDOM (Nov. 29, 2011)).