

From Individual Control to Social Protection: New Paradigms for Privacy Law In The Age of Predictive Analytics

Dennis D. Hirsch

Follow this and additional works at: <https://digitalcommons.law.umaryland.edu/mlr>



Part of the [Law Commons](#)

Recommended Citation

Dennis D. Hirsch, *From Individual Control to Social Protection: New Paradigms for Privacy Law In The Age of Predictive Analytics*, 79 Md. L. Rev. 439 (2020)

Available at: <https://digitalcommons.law.umaryland.edu/mlr/vol79/iss2/4>

This Article is brought to you for free and open access by the Academic Journals at DigitalCommons@UM Carey Law. It has been accepted for inclusion in Maryland Law Review by an authorized editor of DigitalCommons@UM Carey Law. For more information, please contact smccarty@law.umaryland.edu.

**FROM INDIVIDUAL CONTROL TO SOCIAL PROTECTION: NEW
PARADIGMS FOR PRIVACY LAW IN THE AGE OF PREDICTIVE
ANALYTICS**

DENNIS D. HIRSCH*

ABSTRACT

What comes after the control paradigm? For decades, privacy law has sought to provide individuals with notice and choice and so give them control over their personal data. But what happens when this regulatory paradigm breaks down?

Predictive analytics forces us to confront this challenge. Individuals cannot understand how predictive analytics uses their surface data to infer latent, far more sensitive data about them, and so they can no longer make meaningful choices about whether to share their surface data in the first place. Predictive analytics also creates threats (such as harmful bias, manipulation, and procedural unfairness) that go well beyond privacy. Taken together, these two features make it difficult, if not impossible, for traditional privacy law to protect people in the algorithmic economy. If privacy law is to offer meaningful protection, it must shift from a liberalist focus on individual control, to a social protection model in which public authorities set substantive standards that defend people against algorithmic threats.

Leading scholars have recognized the need for such a shift and have proposed ways to achieve it. This Article will argue that, while they move the ball forward, these proposals do not provide an adequate solution. It will propose that the Federal Trade Commission (“FTC”) use its unfairness authority to draw substantive lines between data analytics practices that are socially appropriate and fair, and those that are inappropriate and unfair, and will examine how the FTC would make such determinations. It will

© 2020 Dennis D. Hirsch

* Professor of Law, The Ohio State University Moritz College of Law, and Faculty Director, the OSU Moritz College of Law Program on Data and Governance. Professor of Law, Capital University Law School. I would like to thank Rachel Shonebarger, my research assistant on this article, for her excellent, thoughtful and diligent work. I would also like to thank the members of the Privacy Law Scholar’s Conference, the Midwestern Privacy Law Scholar’s Workshop, and of the Northeast Privacy Law Scholar’s Conference, who workshoped and provided feedback on earlier versions of this paper.

argue that this existing authority, which requires no new legislation, provides a comprehensive and politically legitimate way to create much needed societal boundaries around corporate use of predictive analytics. It will conclude that the FTC could use its unfairness authority to address the threats that the algorithmic economy creates. Were the FTC to do so, that would move us past the liberalist paradigm of privacy law and towards a legal system that can protect people in the age of predictive analytics and artificial intelligence.

ABSTRACT	439
I. INTRODUCTION.....	441
II. THE CONTROL PARADIGM.....	449
III. BEYOND CONTROL	453
A. Predictive Analytics	453
B. The Benefits of Predictive Analytics.....	454
C. Harms from Predictive Analytics	455
D. The Viability of the Control Paradigm.....	459
IV. NEW PARADIGMS FOR PRIVACY LAW AND POLICY	461
A. Information Fiduciaries	464
B. Contextual Integrity	468
C. Use-Based Approach	471
D. Big Data Due Process.....	475
V. PREDICTIVE ANALYTICS AND THE FTC'S UNFAIRNESS AUTHORITY	478
A. Unfairness Authority	479
B. Predictive Analytics and Unfairness Authority: An Example	481
C. Established Public Policies.....	486
D. A Proposal	491
E. Comparing the Various Approaches.....	493
VI. LEGAL AND POLICY QUESTIONS ABOUT THE UNFAIRNESS	
APPROACH.....	496
A. Does the FTC's Section 5 Unfairness Authority Extend to	
Predictive Analytics?	497
B. First Amendment Concerns	502
C. Is Unfairness Authority Too Controversial?	503
VII. CONCLUSION.....	504

I. INTRODUCTION

Some years ago, Target analyzed its customers' past purchasing histories to predict, with great accuracy, which of its current female customers were pregnant. It then showered these women with coupons for baby goods. A high school girl's receipt of the coupons ended up revealing to her father that she was pregnant, a condition that she knew about but had not yet disclosed to her father.¹ Target knew about her pregnancy before her father did, and its actions disclosed this hidden fact to him.

Flooded with tens of thousands of employment applications, Amazon developed an artificial intelligence ("AI") tool that could identify resumes that resembled those of its current employees, and so presumably met Amazon's hiring criteria, and separated them from those that did not. Amazon's existing workforce is disproportionately male. The AI tool accordingly learned to reject applicants whose resumes said that they had gone to all-women colleges or otherwise identified them as female. Fortunately, Amazon caught this bias before it fully implemented the AI tool.²

Cambridge Analytica gained access to a massive trove of Facebook users' "likes." From this information, it was able to infer the personality types of tens of millions of individual Facebook users. It then targeted these individuals with Trump campaign ads designed to appeal, unconsciously, to their specific psychological type³—an approach that *Scientific American* described as a "manipulative strateg[y] . . . designed to bypass one's cognitive defenses."⁴

What do these three examples have in common? Each involves corporate use of predictive analytics, a technological process that analyzes surface

1. Charles Duhigg, *How Companies Learn Your Secrets*, N.Y. TIMES (Feb. 16, 2012), <https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>.

2. Jeffrey Dastin, *Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women*, REUTERS (Oct. 9, 2018), <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.

3. See Kevin Granville, *Facebook and Cambridge Analytica: What You Need to Know as Fallout Widens*, N.Y. TIMES (Mar. 19, 2018), <https://www.nytimes.com/2018/03/19/technology/facebook-cambridge-analytica-explained.html>.

4. Marcello Ienca & Effy Vayena, *Cambridge Analytica and Online Manipulation*, SCI. AM. (Mar. 30, 2018), <https://blogs.scientificamerican.com/observations/cambridge-analytica-and-online-manipulation/>.

data in order to infer and act on the latent information that lies beneath the surface.⁵ Once available only to the government and the largest companies, predictive analytics is today so widely used that it and the technologies based on it—big data analytics, machine learning, and artificial intelligence—have become integral to the information economy.⁶ Companies use these technologies to diagnose diseases, improve medical treatments, design better educational strategies, and make businesses more efficient, to name just a few of predictive analytics' many beneficial applications.⁷ But, as the examples above illustrate, corporate use of predictive analytics also poses risks of privacy invasion (the Target example), discrimination (the Amazon example), and manipulation (the Cambridge Analytica example).⁸ Society needs rules for predictive analytics that will allow this technology to achieve its many benefits, while protecting people against its very real threats.⁹

The existing system of privacy law is not able to protect people against these threats.¹⁰ This is not due to a mistake or an omission. It is due to the very nature of privacy law which was created at a different time and for a different set of technologies. Privacy law seeks to give individuals control over their personal information—to enable them to decide whether to share their personal data, with whom, and for what purpose. But predictive analytics takes surface data and infers latent information from it. This makes it difficult, if not impossible, for people to know what they are really sharing when they agree to disclose their surface data.¹¹ In this way, the rise of predictive analytics undermines the very foundations of existing privacy law.¹²

5. ERIC SIEGEL, PREDICTIVE ANALYTICS 4–5 (2016); U.K. INFO. COMM'R OFFICE, BIG DATA AND DATA PROTECTION 2, 5 (2014), <https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>.

6. Omer Tene & Jules Polonetsky, *Big Data for All: Privacy and User Control in the Age of Analytics*, 11 NW. J. TECH. & INTELL. PROP. 239, 240–43 (2013).

7. VICTOR MAYER-SCHÖNBERGER & KENNETH CUKIER, BIG DATA 1–6 (2013); Dennis Hirsch, *Predictive Analytics Law and Policy: A New Field Emerges*, 14 I/S: J. L. & POL'Y FOR INFO. SOC'Y 1, 4 (2017).

8. Hirsch, *supra* note 7, at 4.

9. WORLD ECONOMIC FORUM, UNLOCKING THE VALUE OF PERSONAL DATA: FROM COLLECTION TO USAGE 3 (2013), http://www3.weforum.org/docs/WEF_IT_UnlockingValuePersonalData_CollectionUsage_Report_2013.pdf (explaining that society needs to govern big data analytics in order to unlock its value).

10. FRED CATE ET AL., DATA PROTECTION PRINCIPLES FOR THE 21ST CENTURY: REVISING THE 1980 OECD GUIDELINES 6–7 (2013), https://www.oii.ox.ac.uk/archive/downloads/publications/Data_Protection_Principles_for_the_21st_Century.pdf; Solon Barocas & Helen Nissenbaum, *Big Data's End Run Around Anonymity and Consent*, in PRIVACY, BIG DATA, AND THE PUBLIC GOOD 44, 45 (Julia Lane et al. eds., 2014); Ira S. Rubinstein, *Big Data: The End of Privacy or a New Beginning?*, 3 INT'L DATA PRIVACY L. 74, 78–79 (2013); Tene & Polonetsky, *supra* note 6, at 242.

11. Neil Richards & Woodrow Hartzog, *Taking Trust Seriously in Privacy Law*, 19 STAN. TECH. L. REV. 431, 446–47 (2016).

12. Woodrow Hartzog, *The Inadequate, Invaluable Fair Information Practices*, 76 MD. L. REV. 952, 953 (2017).

If society is to protect people in the era of predictive analytics, it needs a new strategy grounded in social protection rather than individual control.¹³

To understand where privacy law is today, and where it needs to go, it is helpful to think more broadly about how legal systems evolve. German social theorist Gunther Teubner has explained that Western legal systems begin with a liberalist preference for individual choice over state control.¹⁴ Initially, the state's main role is to support individual choice by enforcing contract and property rights. Over time, however, society increases in complexity to the point that individuals can no longer make meaningful choices about important matters such as the safety of the food or drugs they purchase.¹⁵ In response, the state passes substantive laws, such as food and drug laws, to protect people where they can no longer protect themselves.¹⁶

The law of landlord and tenant clearly illustrates this evolution. Prior to the industrial revolution, courts subscribed to the doctrine of *caveat lessee*—let the lessee beware.¹⁷ Over time, however, simple farm dwellings gave way to tenements and apartment buildings with complex heating, plumbing, and electrical systems.¹⁸ Tenants were no longer able to evaluate the health or safety of such dwellings, and many ended up living in substandard, unhealthy conditions. Courts and legislatures responded with the implied warranty of habitability through which they imposed societal standards designed to ensure safe and healthful housing.¹⁹ A system of law premised on individual choice and control had found itself unable to cope with an increasingly complex set of technologies, and so evolved to encompass societal standards that would better protect people.

That is the same evolutionary shift that privacy law needs to undertake today. Existing U.S. privacy law follows the liberalist model.²⁰ It defines privacy as the ability to control one's personal information.²¹ It then seeks

13. Rubinstein, *supra* note 10, at 74.

14. Gunther Teubner, *Substantive and Reflexive Elements in Modern Law*, 17 *LAW & SOC'Y REV.* 239, 252–54 (1983) (stating that formal law develops into substantive law, which “shifts the focus from autonomy to regulation”).

15. *Id.*

16. *See id.* at 253–54 (“The justification of substantive law is to be found in the perceived need for the collective regulation of economic and social activities to compensate for inadequacies of the market.”).

17. *See* 49 AM. JUR. 2D *Landlord and Tenant* § 449 (2019) (noting that a landlord generally did not have a duty to provide a habitable rental property).

18. *See* *Hilder v. St. Peter*, 478 A.2d 202, 207 (Vt. 1984).

19. *Id.* at 207–08; 15 *WILLISTON ON CONTRACTS* § 48:11 (4th ed. 1990).

20. Neil Richards & Woodrow Hartzog, *Privacy's Trust Gap: A Review*, 126 *YALE L. J.* 1180, 1182 (2017) (stating that ideals of privacy “resonate[] with American ideals of individualism, democracy, and consumerism”).

21. ALAN WESTIN, *PRIVACY AND FREEDOM* 41 (1967); Woodrow Hartzog, *The Case Against Idealising Control*, 4 *EUROPEAN DATA PRO. L. REV.* 423, 423–24 (2018) (“[M]ost people in industry and policy think of privacy and data protection in terms of control. . . . Lawmakers, regulators,

to facilitate such control by requiring companies to notify individuals in advance about the collection and use of their data, give individuals a degree of choice as to whether to allow these data practices, and then employ the information only for the purpose that the individuals authorized.²² This regime's central aim—and so the main goal of U.S. privacy law—is to give individuals control over their personal information and so to give them “privacy” as the law defines it. Professor Paul Schwartz has termed this as the “privacy-control” paradigm of privacy law.²³ The control paradigm is at the core of the Fair Information Practice Principles (“FIPPs”) that have served as a leitmotif for privacy law.²⁴ It is at the heart of most U.S. privacy law statutes.²⁵ And it is reflected in the Federal Trade Commission’s (“FTC”) Section 5 deceptiveness authority which holds companies to the promises they make to individuals about their data practices.²⁶

The control paradigm has worked relatively well for the information collection and processing activities for which it was first designed in the 1970s. However, as happened in food, drug, housing, and so many other areas, increasingly complex social conditions are making it impossible for individuals meaningfully to make choices about the collection and use of their personal information.²⁷ This has been true for some time. Individuals have long struggled to read and assimilate all the privacy notices with which they are confronted.²⁸ But policymakers have remained committed to the

and judges seem to have more or less settled on a notion that the key to privacy generally, and data protection specifically, is control over personal information.”).

22. See DANIEL J. SOLOVE & PAUL M. SCHWARTZ, *INFORMATION PRIVACY LAW* 36–37 (6th ed. 2018).

23. Paul M. Schwartz, *Privacy and Democracy in Cyberspace*, 52 *VAND. L. REV.* 1607, 1613 n.15 (1999).

24. Tene & Polonetsky, *supra* note 6, at 260 (“In the United States, ‘notice and choice’ has been the central axis of privacy regulation for more than a decade.”). The Fair Information Practices have “played a significant role in framing privacy laws in the United States.” Marc Rotenberg, *Fair Information Practices and the Architecture of Privacy (What Larry Doesn’t Get)*, 2001 *STAN. TECH. L. REV.* 1, 15 (2001); see also U.S. DEP’T OF HEALTH, EDUC. & WELFARE, RECORDS, COMPUTERS, AND THE RIGHTS OF CITIZENS viii (1973).

25. See, e.g., Children’s Online Privacy Protection Act, 15 U.S.C. § 6502(b)(1)(A)(ii) (2012) (requiring that children’s website obtain parental consent regarding the collection and disclosure of their children’s personal information); Video Privacy Protection Act, 18 U.S.C. § 2710(b)(2)(B)(ii)(I) (2012) (requiring that a videotape service provider obtain the consent of the consumer for each instance of disclosure); Cable Communications Policy Act, 47 U.S.C. § 551(a)(1) (2012) (requiring cable service providers to notify subscribers of the nature as well as uses of the personal information they collect).

26. Federal Trade Commission Act § 5, 15 U.S.C. § 45(a) (2012).

27. Viktor Mayer-Schönberger & Yann Padova, *Regime Change? Enabling Big Data Through Europe’s New Data Protection Regulation*, 17 *COLUM. SCI. & TECH. L. REV.* 315, 332 (2016).

28. Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 *I/S: J. L. & POL’Y FOR INFO. SOC’Y* 543, 563–64 (2008) (concluding that if individuals read all the policy notices they encountered, the average American would have to spend forty minutes a day reading privacy notices, totaling 244 hours a year); Richards & Hartzog, *supra* note 11, at 444 (declaring individual notice and choice to be “an illusion”).

traditional privacy law approach, trying with mixed success to make it work better, rather than looking in new directions.

That approach is no longer tenable. The rise of predictive analytics has changed the game. Predictive analytics takes massive amounts of data and analyzes them in order to locate correlations and patterns. It then turns these patterns into a profile and applies them prospectively so as to infer additional information and make actionable predictions about the future.²⁹ That is what Target, Amazon, and Cambridge Analytica did in the examples mentioned above. These companies found patterns in data about past customers and used it to make predictions about current customers.³⁰ Today many companies are using this approach to drive their decisionmaking.

The emergence of predictive analytics as a central feature of the digital economy profoundly impacts privacy law's control paradigm.³¹ Previously, individuals who chose to share a particular piece of personal information knew what they were disclosing and could make a meaningful choice about whether or not to do so.³² Today, however, companies can use this surface data to infer additional, latent information that may be far more sensitive than that which the person thought they were sharing.³³ For example, Target used purchase histories to infer pregnancy status; Cambridge Analytica took Facebook "likes" and inferred psychological type. In such a world, individuals cannot understand what information they are really disclosing and, as a consequence, cannot make a meaningful choice about whether or not to share the information in the first place.³⁴

This is not a trivial shift. As will be further explained below,³⁵ in addition to its many benefits, predictive analytics poses significant threats. Companies can violate individual privacy and exploit hidden vulnerabilities. Automated, algorithmic decisionmaking about who should get a loan, a job,

29. See SIEGEL, *supra* note 5, at 4–5.

30. *Cf.*, *id.* at 15 (defining predictive analytics as “[t]echnology that learns from experience (data) to predict the future behavior of individuals in order to drive better decisions.”)

31. Hartzog, *supra* note 12, at 972. In his discussion of algorithmic decisionmaking and the problems it causes for the Fair Information Practices, Hartzog highlights issues that are related, but distinct, from those discussed here, including that “it is very difficult to erase bias from autonomous systems . . . that the cost of these systems [is] not borne equally by all members of society, and . . . that people tend to irrationally trust conclusions reached by computers more than conclusions reached by humans.” *Id.* at 971–72. These concerns, too, are relevant and important.

32. CATE ET AL., *supra* note 10, at 6.

33. *Id.*; Barocas & Nissenbaum, *supra* note 10, at 45–46; Tene & Polonetsky, *supra* note 6, at 240.

34. See Woodrow Hartzog & Evan Selinger, *Big Data in Small Hands*, 66 STAN. L. REV. ONLINE 81, 83 (2013) (stating that predictive inferences make “[e]veryone . . . more susceptible to providing information” because “that [information] gets taken out of its original context” to connect X to Y).

35. See *infra* Sections III.B–C.

insurance, housing, and many other important opportunities, can encode biases and produce a new form of “digital redlining.”³⁶ The black box decision-making processes of predictive analytics are opaque and highly difficult to challenge. Predictive analytics thus threatens people with privacy invasions, manipulation, bias, and procedural unfairness.³⁷

These are real threats. Yet individuals cannot use the tools that privacy law has provided them—notice, choice, and purpose limitation—to protect themselves from these threats. Just as the shift from simple farm dwellings to complex urban ones undermined *caveat lessee* and required lawmakers to develop the implied warranty of habitability, so the move from small to big data puts people at risk and the traditional privacy law paradigm in crisis.³⁸ This shift forces us to ask how the law can protect people in an information economy where, increasingly, individuals cannot protect themselves. It requires us to search for a new regulatory paradigm for the age of predictive analytics.³⁹

This search for a new regulatory paradigm is one of the most compelling inquiries in privacy law, and privacy law scholarship, today. Leading voices recognize that traditional privacy law cannot protect people in a world of predictive analytics; they propose new ways to handle this vital task. Professors Jack Balkin and Jonathan Zittrain and, writing separately, Professors Neil Richards and Woodrow Hartzog, have argued for the expansion of fiduciary duties.⁴⁰ Professors Helen Nissenbaum and Solon Barocas have called for a contextual integrity-based approach.⁴¹ Microsoft’s Craig Mundie and,

36. EXEC. OFFICE OF THE PRESIDENT, BIG DATA: SEIZING OPPORTUNITIES, PRESERVING VALUES 53 (2014), https://obamawhitehouse.archives.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf.

37. See *infra* Section III.C (describing these threats in greater detail); see also Hartzog, *supra* note 12, at 970 (discussing how algorithmic decisionmaking poses risks that include “threats to due process, disparate impact on minority and other vulnerable communities, [and] invasions of privacy and stigmatization” (footnotes omitted)).

38. See, e.g., CHRIS JAY HOOFNAGLE, FEDERAL TRADE COMMISSION: PRIVACY LAW AND POLICY 333 (2016) (“Our regulatory regime, premised on quaint ideas of privacy control . . . is simply inadequate to address the kinds of decision-making and inferential powers that information-intensive industries now possess.”).

39. MAYER-SCHÖNBERGER & CUKIER, *supra* note 7, at 17 (“New principles are needed for the age of big data . . .”); Hartzog, *supra* note 12, at 954; cf. Woodrow Hartzog, *Unfair and Deceptive Robots*, 74 MD. L. REV. 785, 787–88 (2015) (stating that coming age of robots profoundly challenges the existing consumer protection regime).

40. Jack M. Balkin, *Information Fiduciaries and the First Amendment*, 49 U.C. DAVIS L. REV. 1183, 1186 (2016); Richards & Hartzog, *supra* note 11, at 457; Jack M. Balkin & Jonathan Zittrain, *A Grand Bargain to Make Tech Companies Trustworthy*, ATLANTIC (Oct. 3, 2016), <https://www.theatlantic.com/technology/archive/2016/10/information-fiduciary/502346/>.

41. Barocas & Nissenbaum, *supra* note 10, at 47–48.

writing separately, Professor Fred Cate, Professor Victor Mayer-Schönberger, and Peter Cullen have argued for use-based regulation.⁴² Professors Danielle Citron and Frank Pasquale and, writing separately, Professors Kate Crawford and Jason Schultz, have called for importing due process norms into the regulation of big data.⁴³ These authors do not directly refer to one another in their articles, but they are in conversation. Each is proposing a new regulatory paradigm for the age of predictive analytics, one that will protect people where the control paradigm will not.

This Article will describe and evaluate each of these proposals. Building on work in this area⁴⁴ and an essay that this author published in 2015,⁴⁵ this Article will then explore an alternative solution. It will argue that the FTC should use its “unfairness authority” to draw lines between big data practices that are socially appropriate and those that are not, between those that are fair and those that are unfair.⁴⁶

Section 5 of the FTC Act gives the FTC the power to declare and enforce against “unfair or deceptive [business] acts or practices.”⁴⁷ This short phrase gives the FTC two distinct powers: the authority to enforce against business practices that are “deceptive;” and the ability to enforce against business practices that are “unfair.” The FTC has become the nation’s leading privacy

42. Craig Mundie, *Privacy Pragmatism: Focus on Data Use, Not Data Collection*, 93 FOREIGN AFF., Mar./Apr. 2014, at 28, 29 (arguing that we should “shift[] the focus [away] from limiting the collection and retention of data to controlling data . . . the moment when it is used”).

43. Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 23–24 (2014); Kate Crawford & Jason Schultz, *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*, 55 B.C. L. REV. 93, 128 (2014).

44. Citron & Pasquale, *supra* note 43 at 23–24 (calling on the FTC to use its unfairness authority to oversee credit scoring systems); Mark MacCarthy, *New Directions in Privacy: Disclosure, Unfairness and Externalities*, 6 I/S: J. L. & POL’Y FOR INFO. SOC’Y. 425 (2011) (explaining that the FTC could use its unfairness authority to govern corporate data practices).

45. In a 2015 essay, the author briefly explored whether the FTC could use its unfairness authority to draw enforceable lines between fair, and unfair, uses of big data analytics. Dennis D. Hirsch, *That’s Unfair! Or Is it? Big Data, Discrimination and the FTC’s Unfairness Authority*, 103 KY. L.J. 345 (2015). This Article explores in greater depth the FTC’s ability to use its unfairness authority in this way.

46. Other scholars have also identified the FTC’s unfairness authority as a potential mechanism for the governance of emerging technologies. See Hartzog, *supra* note 39, at 811–15 (arguing that the FTC could use its unfairness authority to regulate cognizable consumer harms from robots); Terrell McSweeney, *Psychographics, Predictive Analytics, Artificial Intelligence & Bots: Is the FTC Keeping Pace?* 2 GEO. L. TECH. REV. 514, 522, 525–26 (2018) (explaining that the FTC has cautiously used its unfairness authority to protect privacy and should continue to do so); Richards & Hartzog, *supra* note 11, at 471 (suggesting that the FTC could use its Section 5 authority to regulate unreasonable self-dealing and so enforce fiduciary duties). These works do not focus on predictive analytics *per se* in the way that MacCarthy’s, Citron & Pasquale’s, and this author’s 2015 essay do. MacCarthy, *supra* note 44; Citron & Pasquale, *supra* note 43; Hirsch, *supra* note 45. Taken together, however, this body of work points to the important role that the FTC’s unfairness authority could play with respect to predictive analytics and other information-intensive technologies.

47. 15 U.S.C. § 45(a)(1) (2012).

regulator largely by relying on the first half of this power—its “deceptiveness authority.”⁴⁸ In the algorithmic society, the FTC should put an equal emphasis on the other half of its Section 5 power—its “unfairness authority.”⁴⁹

The FTC’s unfairness authority differs fundamentally from its deceptiveness jurisdiction. Deceptiveness authority seeks to facilitate and enforce individual choices. Unfairness authority does not—it applies in those situations where individuals cannot make meaningful choices.⁵⁰ In those circumstances, it empowers *the FTC* to draw substantive lines between business acts or practices that are socially appropriate and fair and those that are inappropriate and unfair.⁵¹ The move from deceptiveness authority to unfairness authority is thus a shift from a liberalist legal paradigm that seeks to facilitate individual choice and control to one that offers social protection where individuals cannot protect themselves.⁵² Such a shift is needed if the law is to protect individuals in the algorithmic society and allow machine learning, AI, and other forms of predictive analytics to remain socially acceptable.

Part II of this Article will describe the intellectual roots and core features of the control paradigm that currently dominates privacy law. Part III will describe predictive analytics, the risks that it poses to individuals and the broader society, and why control-based regulation is not able effectively to reduce these risks. Part IV will describe and assess the exciting new branch of legal scholarship that has identified this gap and proposed ways to fill it. Part V will explore an alternative: having the FTC use its unfairness authority to draw substantive lines between predictive analytics practices that are fair, and those that are unfair. Part V will then develop a theory of the FTC’s unfairness authority grounded in regulatory history; propose that the FTC employ its unfairness authority to establish rules of the road for predictive analytics and protect people from the threats that it poses; and explain why this approach shares many of the strengths, but avoids many of the weaknesses, of the other main proposals. Part VI will argue that the FTC Act⁵³

48. Daniel Solove & Woodrow Hartzog, *The FTC and the New Common Law of Privacy*, 114 COLUM. L. REV. 583, 628 (2014); see *infra* notes 72–76 and accompanying text.

49. See HOOFNAGLE, *supra* note 38, at 347 (arguing that the FTC should expand its use of its unfairness authority in order to regulate information-intensive industries better). Hoofnagle calls on the FTC to develop a broader theory of how unfairness authority maps onto the advanced information economy. *Id.* This Article attempts to articulate such a theory with respect to algorithmic determinations that negatively impact individuals and society.

50. See 15 U.S.C. § 45(n) (2012) (stating unfairness authority applies to situations where harm to consumers “is not reasonably avoidable by consumers themselves”).

51. HOOFNAGLE, *supra* note 38, at 130–31.

52. MacCarthy, *supra* note 44, at 430 (stating that certain practices are simply unfair even if the information at issue is disclosed to the data subject); see J. Howard Beales, *The FTC’s Use of Unfairness Authority: Its Rise, Fall, and Resurrection*, FTC (May 30, 2003), <https://www.ftc.gov/public-statements/2003/05/ftcs-use-unfairness-authority-its-rise-fall-and-resurrection> (recognizing that the FTC can use its unfairness authority to “attack practices that cause substantial injury,” but could not be addressed under the FTC’s deceptiveness authority).

53. 15 U.S.C. § 45 (2012).

and recent judicial decisions interpreting it provide a sound legal basis for the FTC to use its unfairness authority in this way. This Article will conclude that, in order to protect people in the algorithmic economy, the legal system needs to shift from a liberalist paradigm to one grounded in social protection and that the FTC's unfairness authority provides a practical and politically legitimate way of doing so.

II. THE CONTROL PARADIGM

Many cite Justice Louis Brandeis and Samuel Warren's *The Right to Privacy*⁵⁴ as the main source of U.S. privacy law.⁵⁵ Brandeis and Warren famously argued that those who transgress bounds of appropriateness and so invade others' "right to be let alone" should be held liable in tort.⁵⁶ Their theory is the root of the privacy torts. But it is not the source of U.S. statutory and regulatory privacy law. These are much more grounded in Professor Alan Westin's highly influential 1967 book *Privacy and Freedom*.⁵⁷ It is Westin's book that provides the intellectual underpinnings of modern privacy legislation and regulation.

Professor Westin starts from the idea that humans are social beings who would find living completely unseen by others to be a form of torture.⁵⁸ But they are also individuals who would find it equally painful to be completely transparent to others in their thoughts and actions.⁵⁹ In order to flourish, human beings need to find a balance between these two extremes; different people strike this balance differently. Some are happiest when they share more of themselves with others; some, when they disclose less.⁶⁰ Privacy is the ability to draw this line—to determine which aspects of our personal information we will share with others. "[T]his is the core of the 'right of individual privacy'—the right of the individual to decide for himself, with only extraordinary exceptions in the interests of society, when and on what terms his

54. Samuel Warren & Louis Brandeis, *The Right to Privacy*, 4 HARV. L. REV. 193 (1890).

55. See, e.g., Ben Bratman, *Brandeis and Warren's the Right to Privacy and the Birth of the Right to Privacy*, 69 TENN. L. REV. 623, 624–25 (2002); McSweeney, *supra* note 46, at 516.

56. Warren & Brandeis, *supra* note 54, at 195.

57. WESTIN, *supra* note 21.

58. *Id.* at 40. The stresses of solitary confinement demonstrate this and back up Westin's claim. See Craig Haney & Mona Lynch, *Regulating Prisons of the Future: A Psychological Analysis of Supermax and Solitary Confinement*, 23 N.Y.U. REV. L. & SOC. CHANGE 477, 500 (1997).

59. WESTIN, *supra* note 21, at 41.

60. *Id.* at 40–41.

acts should be revealed to the general public.”⁶¹ Stated a bit differently, privacy is control over the disclosure of one’s personal information.⁶² When one has this control, one has privacy; when one does not, one lacks privacy. Following Professor Schwartz, this Article refers to this as the “control paradigm” of privacy.⁶³

Privacy statutes and regulations that govern the private sector⁶⁴ generally seek to provide individuals with control over their personal information, achieving Professor Westin’s definition of privacy.⁶⁵ They use three core mechanisms to do this: notice, consent, and purpose limitation.⁶⁶ They require companies to *notify* an individual before collecting, using, or sharing that person’s information;⁶⁷ to provide the individual with some degree of choice as to whether or not to *consent* to this data collection and processing;⁶⁸ and to use the information only for the *purposes specified* in the notice to which the person consented.⁶⁹ In theory, these three mainstays of privacy law and policy provide individuals with control over the collection and processing of their personal information and so give them “privacy.”

Notice, consent, and purpose limitation figure importantly in the Fair Information Practice Principles⁷⁰ and are at the heart of most federal privacy

61. *Id.* at 42; *see also* ALAN WESTIN, *PRIVACY AND FREEDOM* xxxi (reissued edition, 2014) (“Most definitions of privacy today agree on the core concept that I presented in 1967: that privacy is the individual’s claim to determine what information about himself or herself should be known to others.”).

62. Hartzog, *supra* note 12, at 973 (identifying this dominant definition of privacy and tracing it back to Westin’s seminal work); Richards & Hartzog, *supra* note 11, at 437.

63. Schwartz, *supra* note 23, at 1664 (referencing the “paradigm of privacy-control”); *accord* Richards & Hartzog, *supra* note 11, at 436 (discussing the “control principle” of privacy).

64. A different legal regime, grounded in the Fourth Amendment, applies to government collection and use of data. *See* Daniel J. Solove, *Digital Dossiers and the Dissipation of Fourth Amendment Privacy*, 75 S. CAL. L. REV. 1083, 1085–86 (2002).

65. Schwartz, *supra* note 23, at 1659 (“From the age of computer mainframes in the 1960s to the current reign of the Internet’s decentralized networks, academics and the law have gravitated towards the idea of privacy as a personal right to control the use of one’s data.”).

66. Barocas & Nissenbaum, *supra* note 10, at 57 (“[I]nformed consent is a natural corollary of the idea that privacy means control over information about oneself.”); Rotenberg, *supra* note 24, at 15 (recognizing that the Fair Information Practices, which center on notice, consent and purpose limitation, have “played a significant role in framing privacy laws in the United States”).

67. *See, e.g.*, Cable Communications Policy Act, 47 U.S.C. § 551(a)(1) (2012) (requiring cable service providers to notify subscribers of the nature as well as uses of the personal information they collect).

68. *See, e.g.*, Children’s Online Privacy Protection Act, 15 U.S.C. § 6502(b)(1)(A)(ii) (requiring that children’s website obtain parental consent regarding the collection and disclosure of their children’s personal information).

69. *See, e.g.*, Video Privacy Protection Act, 18 U.S.C. § 2710(b)(2)(B)(ii)(I) (requiring that a videotape service provider obtain the consent of the consumer for each instance of disclosure). The Act was later amended to allow for consumers to consent to disclosure for a period of two years. *Id.* § 2710(b)(2)(B)(ii)(II).

70. Tene & Polonetsky, *supra* note 6, at 260 (“In the United States, ‘notice and choice’ has been the central axis of privacy regulation for more than a decade.”). The Fair Information Practices

statutes.⁷¹ They also animate the FTC’s “deceptiveness authority.” As was briefly explained above,⁷² Section 5 of the FTC Act authorizes the FTC to identify and enforce against “unfair or deceptive acts or practices” that affect commerce.⁷³ The FTC has successfully asserted that “deceptive behavior” subject to Section 5 enforcement includes those instances in which a company, in a privacy policy or otherwise, makes public representations about how it will or will not collect and use personal information and then turns around and violates this commitment.⁷⁴ The FTC has brought dozens of deception cases against companies that violated their own privacy policies or other public commitments.⁷⁵ These actions have made it the nation’s leading privacy regulator.⁷⁶ The FTC’s deceptiveness complaints support the control paradigm by holding companies to the representations they made to individuals, and on which individuals relied in consenting to company’s processing of their data. Coupled with privacy statutes’ notice and consent requirements, the FTC’s deceptiveness authority establishes the control paradigm as the dominant approach in traditional privacy law and policy.

The control approach depends on two conditions, without which it cannot function. First, individuals must be able to detect when companies collect their data and understand how they will use it. Only then can they make a meaningful choice about whether or not to allow this. Second, companies must be able to know, at the time of collection, the purpose for which they are going to use data. Only then can they notify individuals of this purpose and remain constrained by it. In 1967, when Professor Alan Westin articulated the control approach, most data processing met these conditions.⁷⁷ Companies generally collected personal data in manifest ways and for particular purposes. Individuals could make a meaningful choice about whether to allow this. Even as late as 1980, when the Organization for Economic Co-Operation and Development published its FIPPs,

have “played a significant role in framing privacy laws in the United States.” Rotenberg, *supra* note 24, at 15; *see also* U.S. DEP’T OF HEALTH, EDUC. & WELFARE, *supra* note 24, at viii.

71. *See supra* notes 65–69 and accompanying text.

72. *See supra* notes 47–49 and accompanying text.

73. Federal Trade Commission Act, 15 U.S.C. § 45(a).

74. *See* Complaint at ¶¶ 17, 19, *In re* Snapchat, Inc., No. 132-3078, 2014 WL 1993567, at *3 (F.T.C. May 8, 2014) (asserting that Snapchat’s violation of its privacy policy was subject to the FTC’s deceptiveness authority).

75. Daniel J. Solove & Woodrow Hartzog, *The FTC and the New Common Law of Privacy*, 114 COLUM. L. REV. 583, 600 (2014) (“The FTC has lodged just over 170 privacy-related complaints since 1997, averaging about ten complaints per year.”); *see, e.g.*, *FTC v. Accusearch Inc.*, 570 F.3d 1187, 1190–91 (10th Cir. 2009) (bringing suit against a private company).

76. Solove & Hartzog, *supra* note 75, at 585–86 (“FTC privacy jurisprudence has become the broadest and most influential regulating force on information privacy in the United States—more so than nearly any privacy statute or common law tort.”).

77. “[E]xisting legal frameworks . . . date back from an era of mainframe computers, predating the Internet, mobile, and cloud computing.” Tene & Polonetsky, *supra* note 6, at 241.

[b]usinesses and governments were . . . using personal data in more straightforward ways, often for a single, well-defined purpose Under these circumstances, individuals were more likely to understand the purpose for which their data was being collected and used. And ultimately they could be held accountable for supplying informed consent when given adequate notice.⁷⁸

The digital revolution has made it much more difficult for individuals to achieve this understanding and make these choices.⁷⁹ Today, “[d]ata are generated from online transactions, email, video, images, clickstream, logs, search queries, health records, and social networking interactions; gleaned from increasingly pervasive sensors deployed in infrastructure such as communications networks, electric grids, global positioning satellites, roads and bridges, as well as in homes, clothing, and mobile phones.”⁸⁰ Some of this data creation and collection crosses our awareness, but much does not.⁸¹ Companies seek to give individuals more control over the processing of their data by posting privacy policies that explain their data practices and provide a means (usually opt-out)⁸² for individuals to choose whether to allow them. But the burden of actually reading such notices has become so overwhelming that most people are unable to keep up.⁸³ Policymakers’ main response has been to try to make notice and choice more user-friendly and efficient.⁸⁴

For many, the need to protect privacy meant and continues to mean finding a way to support notice and choice As before, the challenge continues to be perceived as purely operational, as a more urgent need for new and inventive approaches to informing and consenting that truly map onto the states of understanding and assenting that give moral legitimacy to the practices in question.⁸⁵

78. CATE ET AL., *supra* note 10, at 6.

79. Hartzog, *supra* note 12, at 953 (explaining that the shift from information in databases to more advanced information technologies renders inadequate control-based approaches like the Fair Information Practices.)

80. Tene & Polonetsky, *supra* note 6, at 240.

81. For example, companies can harvest data from publicly available records such as government databases or from semi-public data sources such as publicly available Facebook information. *Id.* at 260.

82. SOLOVE & SCHWARTZ, *supra* note 22, at 829 (stating that many privacy policies contain an “opt-out” provision).

83. Hartzog, *supra* note 21, at 428–29; Hartzog, *supra* note 12, at 973, 975; McDonald & Cranor, *supra* note 28, at 563–64 (noting that it would take 244 hours per year for an individual to read through all notices encountered online); Mundie, *supra* note 42, at 30.

84. FTC, PROTECTING CONSUMER PRIVACY IN AN ERA OF RAPID CHANGE: RECOMMENDATIONS FOR BUSINESSES AND POLICYMAKERS 49–50 (2012), <https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf>.

85. Barocas & Nissenbaum, *supra* note 10, at 57–58; Richards & Hartzog, *supra* note 11, at 445 (explaining that “new proposals remain rooted in the Control Illusion”).

The 2018 CONSENT Act⁸⁶ demonstrates this propensity. Proposed by Senators Blumenthal and Markey in response to the Facebook-Cambridge Analytica scandal, this bill would require “edge” providers such as Facebook to notify individuals of their collection and use of sensitive, personally identifiable information and proceed only if the person grants consent.⁸⁷ To be sure, legislators and policymakers have looked to some newer approaches, such as privacy by design.⁸⁸ But proposals such as the CONSENT Act, and the FTC’s continued reliance on its deceptiveness authority,⁸⁹ make clear that the control paradigm remains the dominant policy response to the challenges posed by the digital economy.

III. BEYOND CONTROL

It is time for something more protective. The past dozen years have seen the rapid rise of predictive analytics—a means of analysis that both poses profound risks to individuals and evades the control paradigm’s attempt to protect them. This Part describes predictive analytics and the benefits—and harms—it generates. It then explains why the control paradigm cannot protect people against these threats and so why a new regulatory paradigm is needed to run alongside the control approach and govern predictive analytics.

A. *Predictive Analytics*

Eric Siegel defines predictive analytics as “[t]echnology that learns from experience (data) to predict the future behavior of individuals in order to drive better decisions.”⁹⁰ To unpack this definition, it is helpful to start with an example. Some years ago, a Canadian retailer wanted to identify its credit card holders who were likely to run up bills they could not pay and cut off their credit beforehand.⁹¹ The company had a massive set of data on existing card holders, their purchases, and whether they had defaulted on their payments. It searched this data for a pattern. Did individuals who defaulted on their credit card bills tend to purchase different items than those who diligently paid off their balances? The analysis revealed such a difference.

86. S. 2639, 115th Cong. (2018).

87. *Id.* at § (2)(b)(2)(B); see also Alyson Sandler, *Senate Democrats Propose CONSENT Act*, INSIDE PRIVACY (Apr. 12, 2018), <https://www.insideprivacy.com/united-states/congress/senate-democrats-propose-consent-act/>.

88. FTC, *supra* note 84, at v.

89. McSweeney, *supra* note 46, at 522 (concluding that “[f]or the most part, the FTC continues to rely primarily on its deception authority when policing consumer privacy and the use of consumer data”).

90. SIEGEL, *supra* note 5, at 15.

91. Dana Flavelle, *What the Data Crunchers Know About You*, STAR, (Apr. 23, 2010), https://www.thestar.com/business/tech_news/2010/04/23/what_the_data_crunchers_know_about_you.html.

Those who purchased chrome skull ornaments for the hood of their car or frequented pool halls were more likely to default on their credit card payments.⁹² Those who purchased furniture anti-scuff pads rarely defaulted. The company used this insight to inform its credit decisions.⁹³

This example illustrates well Siegel's definition of predictive analytics. First, the credit card company identified a "target variable"—the attribute it was trying to predict, in this case a low risk of credit card default. Next, it looked to its card holders' purchasing histories. Such "training data" often consists of massive data sets, such as those that the card companies possessed here, that contain the target variable as well as many other data points. It analyzed the training data to find those items that correlated most closely with the target variable. In this case, the strongest association was with furniture anti-scuff pads. Next, the company took that correlation, turned it into a profile, and applied it to individuals not in the original data set to predict whether they possessed the target variable (low risk of credit card default) or not. Where the company saw the "proxy" for low default risk—the purchase of anti-scuff pads—it predicted that the target variable was likely to be present as well. Finally, the company acted on the prediction by marketing its cards to those who had purchased furniture anti-scuff pads. Target variables, correlations, proxies, actionable predictions—these are the core features of predictive analytics.

B. The Benefits of Predictive Analytics

Correlation-based prediction is not new. What is new is the ability to make these predictions with tremendous precision, in real-time, and at scale. In the past dozen years or so, a number of factors have come together to increase dramatically the power of predictive analytics. While it is beyond the scope of this Article to describe them in detail, the main developments include: (1) the rapidly expanding amount of digital information being generated and collected;⁹⁴ (2) the precipitously decreasing cost of collecting and storing that information;⁹⁵ (3) the exponential growth in processing power, as identified in Moore's Law;⁹⁶ and (4) the emergence of "new computational

92. *Id.*

93. *Id.*

94. See SIEGEL, *supra* note 5, at 5; MAYER-SCHONBERGER & CUKIER, *supra* note 39, at 9; WORLD ECONOMIC FORUM, *supra* note 9, at 3 n.1; Rubinstein, *supra* note 10, at 77; Tene & Polonetsky, *supra* note 6, at 240.

95. MAYER-SCHONBERGER & CUKIER, *supra* note 39, at 101 (explaining that, due to technological advances, costs of collection and recordation have dropped significantly, and the cost of digital storage has dropped by roughly fifty percent every two years for the last fifty years); SIEGEL, *supra* note 5, at 6.

96. Rubinstein, *supra* note 10, at 77 (describing "the use of high speed, high-transfer rate computers, coupled with petabytes (i[.].e[.] millions of gigabytes) of storage capacity, resulting in cheap and efficient data processing"); Tene & Polonetsky, *supra* note 6, at 240 (attributing the rise of big data analytics to the "reduced costs of storing information and moving it around in conjunction with

frameworks (such as Apache Hadoop) for storing and analy[z]ing this huge volume of data.”⁹⁷ Together, these technological developments enable data scientists to analyze massive quantities of data in real time to identify a complex string of items that, if found together, predict an individual’s attributes—pregnancy status, likelihood of disease, personality type, etc.—with great precision. An algorithm is defined as “a step-by-step procedure for solving a problem or accomplishing some end”⁹⁸ and is frequently characterized by if/then statements. The string of proxies that, if found together, predict the presence of the target variable are often referred to as an algorithm.

Predictive analytics can be used for great good. It can improve health and save lives, make education more effective, make businesses more productive and efficient, and produce many other improvements and benefits.⁹⁹ Many businesses and economic sectors have adopted predictive analytics in the past dozen years or so, making it increasingly central to our economy and to our lives.¹⁰⁰ Some have gone so far as to identify the “big data economy” as a new economic form that will characterize and inform our society for years to come.¹⁰¹

C. Harms from Predictive Analytics

Unfortunately, predictive analytics can also injure people in significant ways.¹⁰² While many have remarked on this, few have teased out and categorized these threats or assessed whether the control paradigm can successfully address them.¹⁰³ The harms fall into four main categories: privacy invasion, manipulation, bias, and procedural unfairness.¹⁰⁴

increased capacity to instantly analyze heaps of unstructured data using modern experimental methods, observational and longitudinal studies, and large scale simulations”).

97. Rubinstein, *supra* note 10, at 77.

98. *Algorithm*, MERRIAM-WEBSTER, <https://www.merriam-webster.com/dictionary/algorithm> (last visited Feb. 1, 2019); see also Deven R. Desai & Joshua A. Kroll, *Trust But Verify: A Guide to Algorithms and the Law*, 31 HARV. J. L. & TECH. 1, 23 (2017) (defining an algorithm as “a step-by-step process”).

99. Tene & Polonetsky, *supra* note 6, at 241, 247.

100. See Terence Mills, *Eight Ways Big Data and AI Are Changing The Business World*, FORBES (July 31, 2018), <https://www.forbes.com/sites/forbestechcouncil/2018/07/31/eight-ways-big-data-and-ai-are-changing-the-business-world/#5e18915d45b6>.

101. SIEGEL, *supra* note 5, at xxi (explaining that “[m]ore and more, predictive analytics (PA) drives commerce, manufacturing, healthcare, government, and law enforcement”).

102. See Barocas & Nissenbaum, *supra* note 10, at 44; Hirsch, *supra* note 7, at 4.

103. One helpful formulation is from the Future of Privacy Forum. FUTURE OF PRIVACY FORUM, UNFAIRNESS BY ALGORITHM: DISTILLING THE HARMS OF AUTOMATED DECISION-MAKING (2017), <https://fpf.org/wp-content/uploads/2017/12/FPF-Automated-Decision-Making-Harms-and-Mitigation-Charts.pdf>.

104. Hirsch, *supra* note 7, at 4–9 (elaborating on these four categories of harm).

1. *Privacy Invasion*

Predictive analytics takes surface data (like purchase of furniture anti-suff pads) and infers latent data from it (like credit card default risk). This may be why some refer to predictive analytics as “data mining”¹⁰⁵—it unearths and reveals hidden information that lies below the surface. This inferred, latent data can be highly sensitive. For example, as explained above, Target predicted which of its female customers were likely to be pregnant and inadvertently disclosed this fact to one girl’s father.¹⁰⁶ Disclosure of pregnancy status or other sensitive information that the data subject does not wish to reveal is a classic privacy violation.¹⁰⁷

In another example, researchers at Cambridge University were able to predict a person’s gender, sexuality, age, race, and political affiliation “[w]ith remarkable accuracy” based solely on their Facebook likes.¹⁰⁸ As one of those researchers, computational psychologist and big data scientist Michal Kozinski, concluded in a published study:

With just 10 likes, a computer model fundamentally knows you better than a colleague With 70 likes, it knows you better than a friend or roommate; with 150 likes, better than a family member. And with 300 likes, Big Data knows you better than your spouse.¹⁰⁹

Cambridge Analytica (not formally associated with Cambridge University) drew on the Cambridge researchers’ approach to infer 87 million Facebook users’ personality types and target them with political advertisements that they would find hard to resist.¹¹⁰

Purchases at Target, Facebook likes, and the sensitive information that data scientists can infer from them are just the tip of the iceberg. Each of us generates copious data through our online interactions, smart phones, social

105. See Robert Sprague, *Welcome to the Machine: Privacy and Workplace Implications of Predictive Analytics*, 21 RICH. J.L. & TECH. 12, 13 (2015) (referencing “predictive analytics” and “data mining”).

106. See *supra* note 1 and accompanying text.

107. See SOLOVE & SCHWARTZ, *supra* note 22, at 55, 476; Hirsch, *supra* note 7, at 4.

108. Rebecca J. Rosen, *Armed With Facebook ‘Likes’ Alone, Researchers Can Tell Your Race, Gender, and Sexual Orientation*, ATLANTIC (Mar. 12, 2013), <https://www.theatlantic.com/technology/archive/2013/03/armed-with-facebook-likes-alone-researchers-can-tell-your-race-gender-and-sexual-orientation/273963/>.

109. Ben Tinker, *How Facebook “Likes” Predict Race, Religion and Sexual Orientation*, CNN (Apr. 11, 2018), <https://www.cnn.com/2018/04/10/health/facebook-likes-psychographics/index.html>.

110. Cecilia Kang & Sheera Frenkel, *Facebook Says Cambridge Analytica Harvested Data of up to 87 Million Users*, N.Y. TIMES (Apr. 4, 2018), <https://www.nytimes.com/2018/04/04/technology/mark-zuckerberg-testify-congress.html> (stating the number of Facebook users involved); Matthew Rosenberg et al., *How Trump Consultants Exploited the Facebook Data of Millions*, N.Y. TIMES (Mar. 17, 2018), <https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html> (describing how Cambridge Analytica used Facebook likes to infer personality types).

media accounts, internet-connected appliances, etc. From this information, data scientists can infer our most sensitive characteristics and vulnerabilities—our health conditions, sexual preferences, religious beliefs, political commitments—and, if they wish, share these insights with others, all without consulting us, much less getting our consent. Predictive analytics thus poses a profound threat to personal privacy.¹¹¹

2. *Manipulation*

Predictive analytics can reveal people's hidden attributes and vulnerabilities and so can make them vulnerable to manipulation and exploitation. For example, a bad actor might take consumer data, infer from it which particular individuals are likely to be suffering from early stage dementia, and then target them with predatory loans. Or, to take the real world example mentioned above, a company such as Cambridge Analytica might take people's Facebook "likes," use them to infer their personality types, and then target them with political advertisements that they will find hard to resist.¹¹² The difficult questions with respect to manipulation lie in identifying the point at which socially acceptable marketing becomes unacceptable manipulation. An important body of scholarship is beginning to examine this question and suggest ways to draw these lines.¹¹³ Any regulatory system that seeks to reduce the predictive analytics' threats will need to have a method for doing so.

3. *Bias*

Anti-discrimination law distinguishes between claims of disparate treatment and claims of disparate impact. Disparate treatment exists when one intentionally disadvantages others on the basis of their protected class status.¹¹⁴ Disparate impact occurs when a facially neutral standard unintentionally, but meaningfully, disadvantages people on the basis of their protected class status.¹¹⁵ Predictive analytics can produce both types of bias. With

111. Rubinstein, *supra* note 10, at 77.

112. Rosenberg et al., *supra* note 110 (describing how Cambridge Analytica used Facebook likes to infer personality types).

113. See Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995, 998 (2014); Ido Kovaty, *Legally Cognizable Manipulation*, 34 BERKELEY TECH. L.J. 449 (2019); Daniel Susser et al., *Online Manipulation: Hidden Influences in a Digital World*, 4 GEO. TECH. L. REV. 1 (2019); Tal Zarsky, *Privacy and Manipulation in the Digital Age*, 20 THEORETICAL INQ. L. 157 (2019).

114. 45C AM. JUR. 2D *Job Discrimination* § 2395 (2012) (stating that "[u]nder a disparate treatment theory, employment discrimination occurs where an employer has treated a particular person less favorably than others because of a protected trait"); see *Int'l Brotherhood of Teamsters v. United States*, 431 U.S. 324, 335–36 (1977).

115. 9 WEST'S FED. ADMIN. PRAC. 3d § 11856 (2018) (stating that under a disparate impact theory, a plaintiff must show "that the recipient used a facially neutral practice that had a disproportionate impact on a group protected by Title VI"); see *Watson v. Fort Worth Bank and Trust*, 487 U.S. 977, 987–88 (1988); *Int'l Brotherhood of Teamsters*, 431 U.S. at 335 n.15.

respect to disparate treatment, a company could use predictive analytics to infer someone's membership in a protected class (like the class of pregnant women), and then intentionally discriminate against the person on this basis.¹¹⁶

Disparate impact discrimination can occur when existing bias has shaped the training data.¹¹⁷ For example, as mentioned briefly above, Amazon tried to use predictive analytics to evaluate the tens of thousands of resumes it receives.¹¹⁸ It began with the resumes of past applicants and looked for a pattern—which resume attributes correlated closely with being hired at the company, and which did not? The analysis determined that resume items associated with being a woman (like attending an all-female college or being a member of the Women's Chess Club) were associated with unsuccessful applications. The algorithm thus learned to screen out women's resumes. This very likely resulted, not from these women being any less capable, but rather from the documented pro-male bias in the technology industry that had shaped training data—the resumes of successful and unsuccessful job applicants—that the data scientists employed to look for correlations.¹¹⁹

Amazon discovered this problem before implementing the predictive tool.¹²⁰ Had it not done so, the algorithm trained on biased data would have reproduced and perpetuated this bias and done so with the veneer of machine objectivity. As predictive analytics becomes more and more central to eligibility determinations of many types (employment, credit, insurance, school admissions, etc.), harmful bias may be the most significant threat that it creates.

4. Procedural Unfairness

Predictive analytics can produce errors when either the algorithm, or the data to which a company applies the algorithm, is faulty. But human decisionmakers, too, make mistakes. The real problem with algorithmic errors lies in the opacity of the decision-making process and its imperviousness to

116. Such a practice would likely violate employment discrimination laws but would be very hard to detect. *See, e.g.*, The Pregnancy Discrimination Act of 1978, amending Title VII of the Civil Rights Act of 1964, 42 U.S.C. § 2000e (2012) (prohibiting an employer from discriminating against an employee or applicant based on pregnancy status).

117. Solon Barocas & Andrew Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671, 674 (2016); Cynthia Dwork & Deirdre Mulligan, *It's Not Privacy and It's Not Fair*, 66 STAN. L. REV. 35, 37 (2013); Tal Zarsky, *Understanding Discrimination in the Scored Society*, 89 WASH. L. REV. 1375, 1389 (2014).

118. *See supra* note 2 and accompanying text.

119. *Amazon Ditched AI Recruiting Tool That Favored Men for Technical Jobs*, GUARDIAN (Oct. 10, 2018), <https://www.theguardian.com/technology/2018/oct/10/amazon-hiring-ai-gender-bias-recruiting-engine> (noting the pro-male bias in tech companies).

120. *Id.*

challenge. For example, where a predictive algorithm determines that an employee would not succeed in a higher position, and the company accordingly denies him a promotion,¹²¹ the employee in question would have no way to know what data or algorithm resulted in this determination and no way to challenge it.¹²² Such algorithmic determinations are a “black box” as far as the individual is concerned.¹²³ The risk to the individual, then, is that machine-driven decisions deny people the core procedural rights—transparency and the right to be heard—to which they are entitled when others are making important decisions about their lives.

In sum, predictive analytics poses risks of privacy invasion, manipulation, bias (both of the disparate treatment and disparate impact varieties) and procedural unfairness. This is not an exhaustive list. But these are four of the main risks that predictive analytics presents and show that, along with its many benefits, predictive analytics can create some very real harms. The challenge is how to protect people from these threats.

D. The Viability of the Control Paradigm

The control paradigm of notice, consent, and purpose limitation does not sufficiently protect people from the predictive analytics’ harms.¹²⁴ Most people do not understand that predictive analytics takes surface data (e.g., purchase histories) and infers latent information (e.g., pregnancy status) and so cannot make a meaningful choice as to whether to share the surface data in the first place.¹²⁵ The situation is not unlike the renter of an apartment who may be able to see the walls and floors and appliances—the information on the surface—but cannot evaluate the heating and electrical systems.¹²⁶ Just as it makes no sense to apply *caveat lessee* to such a person, it makes no sense to apply the control paradigm to someone who cannot understand what information they are really disclosing.

The second reason the control paradigm cannot protect people today is that the companies that use predictive analytics frequently do not know, at

121. Jeffrey T. Polzer, *Case Study: Should an Algorithm Tell You Who to Promote?*, HARV. BUS. REV. (Feb. 28, 2018), <https://hbr.org/2018/02/case-study-should-an-algorithm-tell-you-who-to-promote>.

122. Rubinstein, *supra* note 10, at 77–78 (“Because decisions based on data mining are largely invisible to their subjects, significant issues arise around access to, and the accuracy and reliability of, the underlying data.”).

123. FRANK PASQUALE, *THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION* 14, 17–18, 218 (2016).

124. McSweeney, *supra* note 46, at 518 (explaining that notice and choice does not adequately protect people against “unanticipated uses of data as inputs for complex algorithms”).

125. CATE ET AL., *supra* note 10, at 6 (stating that given “the proliferation of new information technologies, applications, and data uses, individual consent is rarely exercised as a meaningful choice”); Rubinstein, *supra* note 10, at 78 (“[S]ince users lack knowledge of potential correlations, they cannot knowingly consent to the use of their data for data mining or Big Data analytics.”).

126. *See supra*, notes 17–19 and accompanying text.

the time of collection, the purpose for which they will use a particular piece of data. Companies do not determine the purpose until *after* they analyze the data set and identify the proxies.¹²⁷ It follows that, at the time that they collect data, companies that employ predictive analytics generally do not know, and so cannot provide notice of, the specific purpose for which they will use the data.¹²⁸ “[F]irms that rely on data mining may find it impossible to provide adequate notice for the simple reason that they do not (and cannot) know in advance what they may discover.”¹²⁹

Predictive analytics thus undermines both of the preconditions of the control paradigm. It prevents individuals from understanding what they are really disclosing to companies, and it makes it far more difficult for companies to know and disclose in advance the purpose for which they will use these data. In so doing, predictive analytics prevents people from making meaningful choices about whether to disclose their surface information and so fundamentally undermines the control paradigm of privacy law.

The control approach cannot protect people from the threats that predictive analytics poses. It cannot empower people to protect their privacy because, when people agree to share data, they do not know what they are really revealing. The control approach cannot enable people to protect themselves against manipulation because they often do not know when they are revealing their vulnerabilities. The control approach cannot protect against bias because, while companies may provide notice that they are going to use personal data for analytic purposes, they do not give access to the training data or the algorithm and, even if they did, most people would not know how to assess them.

127. Barocas & Nissenbaum, *supra* note 10, at 60 (explaining that providing notice in the big data context is “challenging, almost by definition, because the value of big data lies in the unexpectedness of the insights that it can reveal”); Rubinstein, *supra* note 10, at 76 (stating that predictive analytics produces newly discovered information that is unintuitive and unpredictable); Tene & Polonetsky, *supra* note 6, at 261 (“Moreover, to be meaningful, consent must be specific to the purpose (or context). Yet by its very nature, big data analysis seeks surprising correlations and produces results that resist prediction.”).

128. A company could, perhaps, notify individuals that it was going to analyze their purchase data so as better to market the company’s goods or to improve services or for one of the other very broad “purposes” that companies sometimes state. But this would tell the purchaser little about how their data was going to be used and how this might affect them. It would not allow the purchaser to protect themselves.

129. Rubinstein, *supra* note 10, at 78; *see also* CATE ET AL., *supra* note 10, at 8 (explaining that “the context in which personal information will be used and the value it will hold are often unclear at the time of collection”); Barocas & Nissenbaum, *supra* note 10, at 60 (big data makes notice and consent difficult because, “[w]ith the best of intentions, holders of large datasets willing to submit them to analyses unguided by explicit hypotheses may discover correlations that they had not sought in advance or anticipated”).

The control paradigm does provide some tools to address procedural unfairness. Both the FIPPs,¹³⁰ and U.S. statutes that implement them,¹³¹ grant individuals the right to access the personal information that a company holds about them, and to correct inaccuracies. Access and correction rights provide some of the transparency and technological due process that people need if they are to contest to algorithmic eligibility determinations that govern their life opportunities (jobs, credit, insurance, etc.). But the FIPPs do not go nearly far enough. They enable an individual only to check and correct their own data, not the entire set of training data on which the algorithm is based. And they do not give them the right to check the algorithm itself to see if it is biased or otherwise faulty. Even with respect to the particular proxy data point that determines a person's eligibility for goods or opportunities, access and correction rights will not be sufficient because individuals, who do not have access to the algorithm, will not know the significance of this data point. Given this, most people will lack an incentive to correct it. Who knew that the purchase of furniture anti-scuff pads was relevant to anything other than stopping one's floors from getting scratched?¹³² No one. And so no one would have an incentive to check on and correct this data point.

For all of these reasons, regulation premised on individual control cannot protect us in a world that is beyond our control.¹³³ “[I]nformed consent is a useful privacy measure in certain circumstances and against certain threats . . . but, against the challenges of big data, consent, by itself, has little traction.”¹³⁴ If the law is to protect people from the threats of predictive analytics—and it must—it needs a new regulatory paradigm with which to do so.¹³⁵

IV. NEW PARADIGMS FOR PRIVACY LAW AND POLICY

A new legal and regulatory paradigm should have two core features. First, it should emphasize social protection, rather than individual control. In

130. Org. for Econ. Cooperation & Dev. [OECD], *Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, ¶13 (2013) (describing the “Individual Participation Principle”), <https://www.oecd.org/internet/ieconomy/oecdguidelinesontheProtectionofPrivacyandTransborderFlowsOfPersonalData.htm>.

131. *See, e.g.*, Cable Communications Policy Act § 631(d), 47 U.S.C. § 551(d) (2012) (requiring cable operators to provide subscribers with “access to all personally identifiable information regarding that subscriber . . . [and a] reasonable opportunity to correct any error in such information”).

132. *See supra* notes 94–96 and accompanying text.

133. CATE ET AL., *supra* note 10, at 6–7; Barocas & Nissenbaum, *supra* note 10, at 45; Rubinstein, *supra* note 10, at 78–79; Tene & Polonetsky, *supra* note 6, at 242.

134. Barocas & Nissenbaum, *supra* note 10, at 58.

135. Rubinstein, *supra* note 10, at 74 (“My contention is that when this advancing wave arrives, it will so overwhelm the core privacy principles of informed choice and data minimization on which the [European Union’s Data Protection Directive 95/46 EC] rests that reform efforts will not be enough. Rather, an adequate response must combine legal reform with the encouragement of new business models . . .”).

other words, it must shift from a liberalist regulatory approach that seeks to facilitate individual choice, to one that empowers public officials to make choices about which predictive analytics practices are safe for individuals and consistent with social values and which are not.¹³⁶ This is the same shift that legislatures and courts made when they moved from *caveat lessee* to the implied warranty of habitability.¹³⁷ In the case of *caveat lessee*, as in the case of privacy, technology evolved to the point that individuals could no longer make meaningful choices and so could no longer protect themselves through their market decisions. In such situations, the liberalist approach to regulation, which seeks to bolster individual choice, must give way to a more interventionist one. Where people cannot protect themselves, society should take steps to protect them.

This shift need not be complete. Where people can still understand what data they are sharing and how they will be used, the liberalist presumption requires that the regulatory system seek to support individual, market choices about these matters.¹³⁸ But where people can no longer make meaningful choices then, by definition, it makes no sense to provide them with the appearance of a choice.¹³⁹ In such situations, legislators and courts should shift from requiring notice and choice to setting substantive rules about which predictive analytics practices are acceptable, and which are not. When it comes to predictive analytics, policy needs to shift from the control paradigm to a “protection paradigm.”

For its second core feature, the new paradigm must be able to protect people against the whole panoply of predictive analytics risks: privacy, manipulation, harmful bias, and procedural unfairness. The big data economy is a new phenomenon that poses new risks and the law must address these threats. If the new paradigm can meet these two criteria—if it can draw substantive lines between those predictive analytics practices that are socially acceptable and those that are not and can do so in a way that protects people against privacy, manipulation, bias, *and* procedural unfairness harms—then it has a chance of protecting people in the algorithmic economy.

Legislators, such as those who proposed the CONSENT Act, seem preoccupied with the control paradigm and are slow to see its limitations or to move beyond it.¹⁴⁰ Academics, however, have been quicker to identify the

136. Barocas & Nissenbaum, *supra* note 10, at 45–46 (explaining that, given the rise of predictive analytics, “procedural approaches cannot replace policies based on substantive moral and political principles”); Hartzog, *supra* note 21, at 431 (“Given the pathologies of mediated choice, people should have a baseline, fundamental level of protection regardless of what they choose.”).

137. *See supra* notes 17–19 and accompanying text.

138. *See* CATE ET AL., *supra* note 10, at 7; Barocas & Nissenbaum, *supra* note 10, at 60 (discussing how notice and consent can still work for smart meters); Hartzog, *supra* note 12, at 954.

139. Barocas & Nissenbaum, *supra* note 10, at 58.

140. *See supra* notes 86–89 and accompanying text.

control paradigm's shortcomings in the era of predictive analytics.¹⁴¹ Several have identified the need for a new regulatory paradigm for reasons similar to those just outlined.¹⁴² They have developed interesting proposals for what this new paradigm should look like.

Professors Jack Balkin and Jonathan Zittrain (in an argument that resonates with the work of Professors Neil Richards and Woodrow Hartzog¹⁴³) propose that certain data-rich companies be treated as “information fiduciaries” and held to duties of care and loyalty in their interactions with their users.¹⁴⁴ Professors Helen Nissenbaum and Solon Barocas maintain that “contextual integrity” theory can provide a conceptual framework for drawing the lines between socially acceptable and unacceptable predictive analytics applications.¹⁴⁵ Professor Fred Cate, Professor Victor Mayer-Schönberger, and Peter Cullen¹⁴⁶ and, writing separately, Craig Mundie, recommend that regulation focus less on the collection and storage of personal data, and more on how companies use it.¹⁴⁷ Professors Danielle Citron and Frank Pasquale and, writing separately, Professors Kate Crawford and Jason Schultz call for a new regulatory framework grounded in due process norms.¹⁴⁸

This body of work is not often spoken of in the same breath. But it should be. Each of these authors recognizes that the control paradigm of privacy law cannot adequately protect people in today's digital economy.¹⁴⁹ Each suggests a way to fill this gap.¹⁵⁰ Their collective search for a new regulatory paradigm to address the threats that predictive analytics poses is

141. Barocas & Nissenbaum, *supra* note 10, at 58 (“While the idea that informed consent *itself* may no longer be a match for challenges posed by big data has been floated by scholars, practitioners, advocates, and even some regulators, such thinking has not entered the mainstream.” (footnote omitted)).

142. See, e.g., CATE ET AL., *supra* note 10, at 6–7; Barocas & Nissenbaum, *supra* note 10, at 45; Hartzog, *supra* note 12, at 977; Richards & Hartzog, *supra* note 11, at 447; Rubinstein, *supra* note 10, at 78–79; Tene & Polonetsky, *supra* note 6, at 242.

143. Richards & Hartzog, *supra* note 11.

144. Balkin, *supra* note 40, at 1186. Professor Balkin and Harvard's Jonathan Zittrain restated and expanded on this argument in a piece in *The Atlantic*. Balkin & Zittrain, *supra* note 40; see also Richards & Hartzog, *supra* note 11, at 457–58 (arguing that privacy law needs fiduciary relationships because they incorporate duties of care and loyalty).

145. Barocas & Nissenbaum, *supra* note 10, at 47.

146. CATE ET AL., *supra* note 10, at 8, 16–17.

147. Mundie, *supra* note 42, at 29.

148. Citron & Pasquale, *supra* note 43 at 18–30; Crawford & Schultz, *supra* note 43, at 109, 120, 128.

149. See CATE ET AL., *supra* note 10, at 6, 8; Balkin, *supra* note 40, at 1200, 1223; Barocas & Nissenbaum, *supra* note 10, at 60; Crawford & Schultz, *supra* note 43, at 108; Mundie, *supra* note 42, at 30; Richards & Hartzog, *supra* note 11, at 434–35; Rubinstein, *supra* note 10, at 78–79.

150. See CATE ET AL., *supra* note 10 (suggesting use-based regulation); Mundie, *supra* note 42 (same); Balkin, *supra* note 40 (proposing fiduciary duties); Richards & Hartzog, *supra* note 11 (same); Barocas & Nissenbaum, *supra* note 10 (recommending contextual integrity); Citron & Pasquale, *supra* note 43 (calling for technological due process); Crawford & Schultz, *supra* note 43 (same).

one of the most exciting areas of privacy law and policy scholarship today. Together with others writing on this topic,¹⁵¹ these scholars are creating a new field focused on the law and policy of predictive analytics.¹⁵² The rest of this Part describes and evaluates the proposals.

A. *Information Fiduciaries*

Privacy law, which seeks to limit the disclosure and use of personal information, and the First Amendment, which seeks to assure its free transmission, can be in tension.¹⁵³ A noted First Amendment scholar, Professor Jack Balkin starts with the idea that any new data protection approach must be able to survive a First Amendment challenge. He frames his influential article, *Information Fiduciaries and the First Amendment*,¹⁵⁴ as an attempt to harmonize privacy regulation with the First Amendment.¹⁵⁵

Balkin focuses on an important exception to the First Amendment: fiduciary duties of loyalty and of care.¹⁵⁶ Such duties can require a lawyer or doctor, for example, to use a client's or patient's personal data only in ways that benefit that person and to obtain the client's or patient's consent before doing so.¹⁵⁷ Such restrictions do not offend the First Amendment.

Balkin employs this exception to craft a policy proposal that could, consistent with the First Amendment, put substantive constraints on some commercial data practices. He points out that social media platforms, search engines, and certain other online service providers resemble fiduciaries.¹⁵⁸ Much like lawyers or doctors, these service providers receive sensitive personal information they are expected to treat with care and use to further the data subject's interests. Balkin calls online service providers "information

151. Hartzog, *supra* note 12, at 956 (arguing for greater regulatory attention to the design of information technology); Rubinstein, *supra* note 10 (advocating for sharing the benefits of predictive analytics with individual users); Tene & Polonetsky, *supra* note 6 (same). See generally Audio Recording: Symposium on Predictive Analytics Law and Policy: Mapping the Terrain, held by I/S: Journal of Law and Policy for the Information Society (Mar. 24, 2017), <https://moritz-law.osu.edu/briefing-room/multimedia/2017-is-symposium/>.

152. See Hirsch, *supra* note 7, at 9.

153. Balkin, *supra* note 40, at 1194 (stating that "the First Amendment puts rather strict limits on how government might regulate companies . . . that collect large amounts of information about end-users and then analyze, use, distribute, and sell that information to make profits, or to gain business or political advantages").

154. *Id.* at 1183.

155. *Id.* at 1186 ("This essay attempts to make these two commitments cohere—to show how protections of personal privacy in the digital age can co-exist with rights to collect, analyze, and distribute information that are protected under the First Amendment.").

156. *Id.*

157. *Id.* at 1208.

158. *Id.* at 1125, 1129.

fiduciaries” and argues that the law could, consistent with the First Amendment, impose on them duties of care and loyalty.¹⁵⁹ These duties would extend to the information fiduciary’s use of predictive analytics to make decisions about people.¹⁶⁰ In a separate, significant article, Professors Neil Richards and Woodrow Hartzog also look to fiduciary law as a model for privacy regulation.¹⁶¹ This analysis will focus on Professor Balkin’s work, but much of it would apply to Professors Richards and Hartzog’s important proposal as well.

Fiduciary duties are substantive obligations that go well beyond notice, choice, and purpose limitation. They seek, not to give individuals control over their personal information, but to impose substantive limits on the fiduciary’s actions with respect to this data. In this way, Balkin’s proposal satisfies the first requirement of the new regulatory paradigm—that it create substantive rules that draw the line between algorithmic activities that are socially acceptable, and those that are not. The “duty of care” obliges fiduciaries to “take care to act competently and diligently so as not to harm the interests of the principal, beneficiary, or client.”¹⁶² The “duty of loyalty” requires them to “keep their clients’ interests in mind” and, where the fiduciary’s interests might conflict with the clients’, always to “act in their clients’ interests.”¹⁶³ Balkin recommends that Congress pass legislation imposing these obligations on online service providers. This, he maintains, would create substantive limits on commercial data practices and help to protect individuals in today’s digital economy.

In a subsequent article published in *The Atlantic* and titled *A Grand Bargain to Make Tech Companies Trustworthy*, Professor Balkin and Professor Jonathan Zittrain call for a federal statute that would give tech companies the option of taking on these fiduciary duties with respect to their users.¹⁶⁴ Such companies would agree to “the duty to use personal data in ways that don’t betray end users and harm them.”¹⁶⁵ In exchange, the statute would pre-empt the “patchwork of state and local laws about online privacy” that such companies find burdensome and difficult to comply with.¹⁶⁶ This “grand bargain”

159. *Id.* at 1209 (“An information fiduciary is a person or business who, because of their relationship with another, has taken on special duties with respect to the information they obtain in the course of the relationship.”).

160. *Id.* at 1232.

161. Richards & Hartzog, *supra* note 11, at 457 (stating that fiduciary law is aimed at protecting “against the exploitation of a vulnerability created by trust in another”).

162. Balkin, *supra* note, 40, at 1208.

163. *Id.* at 1207–08.

164. *Id.*

165. Balkin & Zittrain, *supra* note 40.

166. *Id.* (noting that complying with varying state legislation has “become sufficiently burdensome”).

could create business support for such legislation and so facilitate its passage.¹⁶⁷

Professor Balkin's information fiduciary idea proposal contains the two required features. It imposes substantive obligations—the duty of care and the duty of loyalty—on online service providers. And it reaches beyond privacy harms to encompass predictive analytics' other risks. Fiduciary duties “include duties not to use information obtained in the course of the relationship in ways that harm or undermine the principal, patient, or client, or create conflicts of interest with the principal, patient, or client.”¹⁶⁸ This broad limitation should prevent, not just privacy injuries, but also manipulation, bias, and procedural unfairness, each of which could harm or undermine the individual. By casting the obligations as fiduciary duties, Balkin's proposal obviates any potential First Amendment obstacle to substantive limits of this type.¹⁶⁹

For all of its advantages, Balkin's proposal will not protect people sufficiently from the risks of predictive analytics. To begin with, the proposal's scope is too narrow. It covers only those “information fiduciaries” who “[b]ecause of their special power over others and their special relationships to others, . . . have special duties to act in ways that do not harm the interests of the people whose information they collect, analyze, use, sell, and distribute.”¹⁷⁰ It does not cover the large number of digital businesses that have no such direct relationships with individuals. This is an important gap. Increasingly, companies use predictive analytics to make eligibility determinations¹⁷¹ for marketing offers, jobs, loans, insurance, housing, and other vital goods and life opportunities.¹⁷² This aspect of the algorithmic economy—what Professors Citron and Pasquale have called “technology-driven adjudication”¹⁷³—constitutes one of the most significant forums in which privacy violations, bias, procedural unfairness, and manipulation can rear their ugly heads. A new paradigm for regulating predictive analytics must cover this situation and prevent these harms.

Professors Balkin and Zittrain's information fiduciaries model would not. Generally, the companies making these eligibility determinations do not have a relationship with the individual at the time the person applies for the loan, job, insurance, housing, or other good—the critical moment at which predictive analytics can work its harm. Without such a relationship, these companies will not meet Balkin's definition of information fiduciary and will

167. *Id.*

168. Balkin, *supra* note 40, at 1208.

169. *Id.* at 1209–10.

170. *Id.* at 1186.

171. *See* Citron & Pasquale, *supra* note 43, at 2–3.

172. *See id.* at 4.

173. *Id.* at 19.

not be covered by his proposed law.¹⁷⁴ Professor Balkin's proposal thus leaves out the very companies that can do the most damage: those that employ predictive analytics to make eligibility determinations. Balkin recognizes this limitation:

[T]he analysis I have offered in this essay can only take us so far. The concept of information fiduciaries presented here focuses on the violations of special relationships between companies and the people whose information they collect, collate, and use. . . . But in the Algorithmic Society, companies will purchase and use lots of data that is not so encumbered, and they will use it to affect the lives of countless people who are not their clients or end-users.

At this point, we can no longer rely on the notion of special fiduciary relationships between individuals and companies to regulate the use and abuse of data.¹⁷⁵

Even if information fiduciary theory were able to overcome this scope limitation and establish a relationship between the company and the individual, the duties of care and loyalty that the theory prescribes would not be appropriate for large swaths of the algorithmic economy.¹⁷⁶ Consider a company that utilizes predictive analytics to determine which of its existing employees to interview for a promotion to a higher-level job. The algorithm evaluates Ed Employee, predicts that he would not perform well in the new position, and so the company does not consider him for it.¹⁷⁷ Under information fiduciary theory, the company would just have violated the duty of loyalty.¹⁷⁸ That is, it would have put its own interests ahead of the data subject's interests.¹⁷⁹ But prioritizing the employer's needs is just what society would expect and want the employer to do when selecting someone for a promotion.

174. Balkin, *supra* note 40, at 1232–33 (recognizing that, “when algorithms use data about *other* people,” there is no violation of a fiduciary relationship).

175. *Id.* at 1233.

176. Professors Richards and Hartzog take a slightly more nuanced approach. While they call for making the duty of loyalty “a foundational concept in privacy law,” Richards & Hartzog, *supra* note 11, at 468, they define the duty more narrowly and argue that the law should regulate “unreasonable self-dealing.” *Id.* at 471. However, they do not explain how to draw the line between reasonable, and unreasonable, self-dealing.

177. This is the field of “people analytics” and is much used in business today. Carla Arellano et al., *Using People Analytics to Drive Business Performance: A Case Study*, MCKINSEY & CO. (July 2017), <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/using-people-analytics-to-drive-business-performance-a-case-study>.

178. Balkin, *supra* note 40, at 1209 (explaining that information fiduciaries have a duty “to use the information they obtain about their clients for the client’s benefit and not to use the information to the client’s disadvantage”).

179. *Id.*

This example highlights an important limitation on the information fiduciary approach: Fiduciary duties are not designed to apply to all companies. They are designed for those commercial actors, such as lawyers, doctors, or accountants, whose very function requires them to put their clients' interests ahead of their own. They do not apply, and should not apply, to the great majority of commercial actors who, with limited exceptions,¹⁸⁰ are *expected* to pursue their own self-interest. Society requires such companies to deal fairly and honestly with others. But it does not demand that they behave like a fiduciary. This further limits the usefulness of Professor Balkin's approach when it comes to governing the algorithmic economy. Cameron Kerry's interesting proposal for a "Golden Rule of Privacy," under which those entrusted with another's personal information would have "the obligation to act in the interests of the beneficiaries and to avoid self-dealing," is similar in many ways to Professor Balkin's fiduciary duties and suffers from this same limitation.¹⁸¹

Another issue with the information fiduciary model is that beneficiaries can often consent to the disclosure of their information by a fiduciary who might otherwise have a duty to keep it confidential. Thus, the information fiduciary approach does not fully free us from the consent problem.

The final problem with Professor Balkin's proposal is that it would require congressional action. While Balkin and Zittrain make an interesting case for why members of Congress and tech companies should support such a statute, getting legislation through Congress remains an uphill climb. Analytics-based eligibility determinations are being made today that will have lasting effects. People need regulatory protections now.

B. Contextual Integrity

In their illuminating book chapter, *Big Data's End Run Around Anonymity and Consent*, Professors Helen Nissenbaum and Solon Barocas offer an alternative regulatory framework for predictive analytics,¹⁸² one grounded in Nissenbaum's highly influential theory of contextual integrity.¹⁸³ Like Professor Balkin in his *Information Fiduciaries* piece¹⁸⁴ and this Article,¹⁸⁵ Professors Nissenbaum and Barocas recognize that notice, consent, and the other elements of the control paradigm do not protect people sufficiently

180. For example, they cannot discriminate against protected classes even if they believe that doing so would further their own interests.

181. Cameron F. Kerry, *Why Protecting Privacy Is a Losing Game Today—and How to Change the Game*, BROOKINGS (July 12, 2018), <https://www.brookings.edu/research/why-protecting-privacy-is-a-losing-game-today-and-how-to-change-the-game/>.

182. Barocas & Nissenbaum, *supra* note 10.

183. HELEN NISSENBAUM, *PRIVACY IN CONTEXT: TECHNOLOGY, POLICY, AND THE INTEGRITY OF SOCIAL LIFE* (2010).

184. *See supra* notes 154–167 and accompanying text.

185. *See supra* Part II.

from the harms of predictive analytics.¹⁸⁶ If individual choice is not to establish limits on this type of data processing, what should take its place?

Nissenbaum and Barocas develop a two-part answer. First, the informational norms that surround the collection and use of particular pieces of information should limit business activities with respect to this data. People generally share information in a particular context (e.g., a professional office or a friendship). Social norms specific to that setting govern these data flows.¹⁸⁷ Business use of predictive analytics that accords with these context-specific norms is presumptively appropriate; that which does not, is presumptively inappropriate.¹⁸⁸ In this way, social norms provide a basis for drawing substantive lines between acceptable and unacceptable applications of predictive analytics.

This cannot be the end of the inquiry, however. Norms change in response to evolving social conditions and technologies.¹⁸⁹ In their day, Brandeis and Warren were outraged by the emergence of the “snap” camera that could take photos in an instant.¹⁹⁰ Today, we accept such cameras without a second thought. A legal regime that outlaws all data practices inconsistent with existing informational norms would impede the process by which new technologies, including those highly beneficial to humanity, are introduced and norms change. Professors Nissenbaum and Barocas’s commitment to informational norms thus creates a problem for their theory.

To save the contextual integrity theory, they must identify a way to distinguish norm-breaking data practices that are legitimate and acceptable, from those that are not.¹⁹¹ Nissenbaum and Barocas articulate a way to draw this line. Data practices that transgress informational norms are permissible where they are “more effective in promoting interests, general moral and political values, and context-specific ends, purposes, and values” such as “fairness, justice, freedom, autonomy, welfare, and others more specific to the context in question,” than those practices that comply with existing informational norms.¹⁹² Data practices that are not more effective at doing so are impermissible.

186. Barocas & Nissenbaum, *supra* note 10, at 45 (explaining that “[i]n practice, . . . anonymity and consent have proven elusive, as time and again critics have revealed fundamental problems in implementing both”).

187. *Id.* at 47.

188. “Entrenched norms” are the default. *Id.* at 48.

189. *Id.* at 47.

190. See Warren & Brandeis, *supra* note 54, at 195 (“Instantaneous photographs and newspaper enterprise have invaded the sacred precincts of private and domestic life . . .”).

191. Barocas & Nissenbaum, *supra* note 10, at 47 (“To explain why such disruptions are morally problematic—or rather to distinguish between those that are and are not—a norm-based account of privacy, such as contextual integrity, must offer a basis for drawing such distinctions.”).

192. *Id.* at 48. “For the theory of contextual integrity, the touchstones of moral legitimacy include interests and general moral and political values (and associated rights), commonly cited in

Nissenbaum and Barocas's approach thus requires two steps. First, ask whether a particular data flow is consistent with context-specific, existing informational norms. If so, the data flow is presumptively legitimate; if not, it is presumptively illegitimate. Second, ask whether the data practice that transgresses existing informational norms promotes moral, political and context-specific values such as "fairness, justice, freedom, autonomy, welfare, and others more specific to the context in question" better than existing, norm-compliant practices.¹⁹³ If so, this can rebut the presumption and render the new data practice legitimate and appropriate, notwithstanding the fact that it breaks existing informational norms. Professors Nissenbaum and Barocas phrase it in this way: "a heuristic supported by contextual integrity sets entrenched norms as default but allows that if novel practices are more effective in promoting interests, general moral and political values, and context-specific ends, purposes, and values, they should be favored over the status quo."¹⁹⁴

Contextual integrity theory does important work with respect to the regulation of predictive analytics. Rather than rely on the illusion of individual control, it draws a substantive line between data practices that are appropriate and those that are not. It further highlights existing informational norms as important points of reference for making this determination.

But it is less helpful when it comes to deciding whether predictive analytics practices that break with social norms (which will likely be most such operations, given the newness of the field and the premium that it places on innovation) are nonetheless appropriate and acceptable. To determine whether such practices are acceptable, the theory would require one to assess whether the norm-breaking analytic practice is better able to promote "interests, general moral and political values, and context-specific ends, purposes, and values" such as "fairness, justice, freedom, autonomy, welfare, and others more specific to the context in question," than a norm-compliant one would be.¹⁹⁵ This test is so vague as to be almost unworkable. Which general and context-specific interests and values is one to consider? Professors Nissenbaum and Barocas point us to certain values—fairness, justice, freedom, autonomy, welfare. But each of these grand concepts carries with it a lengthy intellectual history and is highly contestable. And who is to say that Nissenbaum and Barocas have even arrived at the right list? Might there be other values capable of overriding existing informational norms such as "family" or "democracy" or any number of other highly important values?

accounts of privacy. Beyond these, however, a further distinctive set of considerations are context-specific ends, purposes, and values." *Id.* at 47–48.

193. *Id.* at 48.

194. *Id.*

195. *Id.*

Assuming that the values themselves were clear and uncontested, the decisionmaker would still face the problem of deciding whether the data practice in question would further these values better than alternative practices that are consistent with existing informational norms. This would require the decisionmaker to define the operative values (“fairness, justice, freedom, autonomy,” etc.) and then make a judgment about how they apply to a given context. Each of these steps is difficult and subjective. How is one to define clearly “freedom” or “justice” or “autonomy”? And how is one to know, with sufficient certainty, whether a given analytic innovation will further these worthy goals? Professors Nissenbaum and Barocas themselves acknowledge that “since the world is a messy place, rife with conflict and uncertainty, it is usually on the basis of partial knowledge only that we seek to optimize on these factors.”¹⁹⁶ The contextual integrity approach to predictive analytics suffers from excessive vagueness and ambiguity.

It also pays insufficient attention to political legitimacy. Whose concept of freedom, or justice, or autonomy is to govern here? The regulator’s? A philosopher’s (which one)? The company’s own vision? Constitutional democracies have a way of settling contested values questions. To state the obvious, they hold elections and then let the elected representatives and those to whom they delegate power choose among competing values.¹⁹⁷ But Nissenbaum and Barocas do not root the values in these politically legitimate decisionmakers. Who, then, is to be the arbiter? And how are this entity’s decisions to be squared with our democratic and Constitutional commitments? Finally, as with Professor Balkin’s information fiduciary approach, current law provides no broad right to contextual integrity and our gridlocked Congress would have to enact such a limit. This is a major practical hurdle.

C. Use-Based Approach

Two works offer an even more radical departure from traditional notice-and-choice-based privacy regulation. Professor Fred Cate, Professor Victor Mayer-Schönberger, and Peter Cullen’s white paper, *Data Protection Principles for the 21st Century*,¹⁹⁸ and Craig Mundie’s article *Privacy Pragmatism*,¹⁹⁹ each begin with the idea that, under traditional privacy regulation, people are to receive notice and make a choice *at the moment of data collection*.²⁰⁰ As they see it, this is the wrong place in the data lifecycle on which

196. *Id.*

197. Paul Nemitz, *Constitutional Democracy and Technology in the Age of Artificial Intelligence*, 376 PHIL. TRANSACTIONS A, Aug. 14, 2018, at 1.

198. CATE ET AL., *supra* note 10. Cate, Cullen, and Mayer-Schönberger’s paper reports the conclusions of a Microsoft-sponsored working group, of which the three authors were members, that was organized by the Oxford Internet Institute. *Id.* at 11–12 (describing the working group and naming its members).

199. Mundie, *supra* note 42.

200. *Id.* at 29.

to focus, for three reasons. First, in an era of ubiquitous data collection and analytics-driven inferences, people do not have the time or the knowledge to make intelligent choices about whether to allow such collection.²⁰¹ Second, limits on collection can prevent companies from assembling the data that they need to make socially beneficial predictions, such as those that improve health care.²⁰² Finally, when data processing harms people it generally does so at the point that the processor *uses* the data, not when it collects it.²⁰³

These authors accordingly argue that the privacy law system should shift its focus from collection and storage, to data use.²⁰⁴ Privacy law should allow companies to collect and store personal data without constraint. But the law should limit the ways in which companies can utilize the data that they have collected. “The time has come for a new approach: shifting the focus from limiting the collection and retention of data to controlling data at the most important point—the moment when it is used.”²⁰⁵ The President’s Council of Advisors on Science and Technology, in a high-profile 2014 report on big data and privacy, relied heavily on Mundie’s article and endorsed the use-based approach.²⁰⁶

While the two papers agree on the superiority of a use-based approach, they differ on how to operationalize it. Professor Cate, Professor Mayer-Schönberger and Cullen argue that the permissibility of a given use should depend on whether its benefits likely outweigh its harms, taking into consideration the measures put into place to reduce those harms.²⁰⁷ Recognizing that such benefit-harm balancing must necessarily require difficult, values-based, “context-specific risk assessment[s],” the authors call on the legislature to pass laws that would “determine clearly how harms and benefits are to be evaluated.”²⁰⁸ The authors further single out data uses that affect a

201. CATE ET AL., *supra* note 10, at 6–7; Mundie, *supra* note 42, at 30–32.

202. Mundie, *supra* note 42, at 33 (discussing Kaiser Permanente’s use of big data to find a link between expectant mothers’ use of anti-depressant drugs and autism spectrum disorders among their children).

203. *Id.* at 33–34.

204. CATE ET AL., *supra* note 10, at 8; Mundie, *supra* note 42, at 29. Fred Cate articulated the main contours of this argument as early as 2006. See Fred Cate, *The Failure of Fair Information Practice Principles*, in CONSUMER PROTECTION IN THE INFORMATION ECONOMY 341 (Jane K. Winn ed., 2006). Cate’s 2006 proposal overlaps substantially with the 2014 white paper that he co-authored. CATE ET AL., *supra* note 10.

205. Mundie, *supra* note 42, at 29; see CATE ET AL., *supra* note 10, at 11 (stating that a priority for modernizing the central premises of privacy law is to “[r]educe the focus on data collection and the attending notice and consent requirements, and focus more on a practical assessment of the benefits and risks associated with data uses”).

206. PRESIDENT’S COUNCIL OF ADVISORS ON SCI. & TECH., EXEC. OFFICE OF THE PRESIDENT, BIG DATA AND PRIVACY: A TECHNOLOGICAL PERSPECTIVE 41 (2014), https://bigdata.tawg.nist.gov/pdf/pcast_big_data_and_privacy_-_may_2014.pdf.

207. CATE ET AL., *supra* note 10, at 17.

208. *Id.* at 17–18.

person's "education, employment, physical or mental health, financial position, or legally protected rights."²⁰⁹ They maintain that companies should notify individuals of such uses, allow them to access and correct the data about themselves, and permit them to challenge "the processing and accuracy" of that data.²¹⁰ This appears to allow individuals to contest the company's benefit-cost assessment and, ultimately, to challenge it in court.

Mundie relies, not on benefit-cost balancing, but on a tech-savvy form of individual choice to govern data use.²¹¹ He proposes the development of software that would "wrap" an individual's personal information at the point of origin with meta-data that expresses the individual's preferences regarding use.²¹² The wrapper would describe the content of the data within it but would not allow the processor to access the data itself until it agreed to abide by the terms of use that the data subject had specified.²¹³ Individuals who believed that their data had been misused could file a complaint with a regulator who, if it found such misuse, would issue significant penalties.²¹⁴ Such a system, Mundie argues, would allow individuals to "make more informed decisions" about whether to allow companies to use their data²¹⁵ and so protect themselves against harmful uses.

The use-based approach to predictive analytics has a number of strengths. It correctly recognizes the limits of notice and choice regulation, particularly as applied to predictive analytics. It shines a useful light on data use and on the harms that can occur at this point of the data lifecycle. And Cate, Cullen, and Mayer-Schönberger's benefit-cost component has some appeal in an area that promises both important benefits and significant risks.²¹⁶

The use-based approach also has significant problems. To begin with, companies can harm individuals not only by using personal data, but also by collecting and storing it. For example, televisions that continuously track and report what consumers are watching²¹⁷ or the always-on, in-home voice as-

209. *Id.* at 18.

210. *Id.* at 18–19.

211. Mundie, *supra* note 42, at 34.

212. *Id.*

213. *Id.*

214. *Id.* at 36.

215. *Id.* at 34.

216. *Cf.* Tene & Polonetsky, *supra* note 6, at 241 ("The extraordinary societal benefits of big data—including breakthroughs in medicine, data security, and energy use—must be reconciled with increased risks to individuals' privacy.").

217. *See* Complaint for Permanent Injunction and Other Equitable Relief at 4, *FTC v. Vizio, Inc.*, No. 2:17-cv-00758, 2017 WL 7000553 (D.N.J. Feb. 6, 2017) (alleging the Vizio TVs "continuously track what consumers are watching, and transmit that information to Defendants").

sistant that inadvertently recorded and transmitted a couple's private conversation²¹⁸ affect people simply by collecting data. The use-based approach would give companies free reign to collect personal information through these and other internet-enabled surveillance devices. Left unchecked, this could create an always-on surveillance environment in which our most personal moments would be recorded and stored. This would chill our free expression and push us towards conformity.²¹⁹ The advocates of the use-based approach are too quick to give companies free reign to conduct surveillance of this type. Harm can occur at any point in the data lifecycle—from collection, to storage, to use—and an approach that recognizes this will be more protective than one focused solely on use.

Mundie's model of use regulation has a further problem. The "wrapper" around the data is supposed to specify all allowed uses. But there are far too many such uses for a person to understand and determine in advance. Much like the notice and choice paradigm before it, Mundie's use-based model ends up putting an overwhelming burden on the individual. This shortcoming becomes even more profound when applied to predictive analytics. Companies that employ algorithmic decisionmaking often do not know in advance how they will use data that they collect.²²⁰ They discover this later through the search for correlations. Mundie's approach assumes that data processors *will* know the uses to which they want to put an individual's data at the moment they collect it, for that is the only way in which a company could, at that moment, agree to the data subject's terms of use. Mundie's use-based approach will not work for predictive analytics.

Cate, Cullen, and Mayer-Schönberger's version has more potential. Its benefit-cost approach puts the onus on a regulator, not the individual, to make choices about potential uses. It is also framed broadly enough to apply, not just to privacy invasions, but also to bias, exploitation, and manipulation. Coupled with the procedural protections that they recommend for especially risky data processing, Cate, Cullen, and Mayer-Schönberger's model begins to define a comprehensive approach that could allow people to challenge algorithmic determinations and provide a substantive basis for evaluating these claims.

But Cate, Cullen, and Mayer-Schönberger leave far too much unsaid. Where the benefits and costs are not readily reducible to dollars and cents, by what metric should they be balanced? Their proposal says only that "harms should be permitted with protections in place appropriate to the risk

218. Niraj Chokshi, *Is Alexa Listening? Amazon Echo Sent Out Recording of Couple's Conversation*, N.Y. TIMES (May 25, 2018), <https://www.nytimes.com/2018/05/25/business/amazon-alexa-conversation-shared-echo.html>.

219. Helen Nissenbaum, *Deregulating Collection: Must Privacy Give Way to Use Regulation?* 21 (May 1, 2017) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3092282.

220. *See supra* notes 127–129 and accompanying text.

and degree of harm,” and that each nation should pass legislation clarifying how this will work.²²¹ But the proposal provides little guidance on how to fashion the statutory test. They lead us towards a solution but leave us there without essential pieces of the puzzle.²²²

D. Big Data Due Process

Two important articles published within months of one another—Professors Kate Crawford and Jason Schultz’s *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*,²²³ and Professors Danielle Keats Citron and Frank Pasquale’s *The Scored Society: Due Process for Automated Predictions*²²⁴—focus on procedural protections for those who are the subjects of, and affected by, analytic processes. Drawing on Professor Citron’s path-breaking article *Technological Due Process*,²²⁵ these articles assert that when companies engage in algorithmic adjudications,²²⁶ they are wielding a quasi-governmental authority that can profoundly affect a person’s life opportunities.²²⁷ The “underlying values of due process,”²²⁸ especially transparency and the right to be heard,²²⁹ should accordingly apply.

With regard to transparency, the authors maintain that companies should be required to disclose to data subjects and regulators “the issues that were predicted, and ideally, the data considered and the methodology employed.”²³⁰ Companies should further create and store an “audit trail that records the basis of the predictive decisions, both in terms of the data used and the algorithm employed”²³¹ and give individuals and regulators the opportunity to access this audit trail on demand.²³² If providing this information to individuals posed too great a threat to trade secrets or intellectual property interests, the company could provide it to an “independent third part[y]” that

221. CATE ET AL., *supra* note 10, at 17.

222. Cate, Cullen and Mayer-Schönberger would provide such individuals with notice of the determination and an opportunity to challenge its accuracy. *Id.* at 18. But that does not answer the question of whether the data use, even if accurate, is not socially acceptable and fair.

223. Crawford & Schultz, *supra* note 43.

224. Citron & Pasquale, *supra* note 43.

225. Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249 (2008). Each of the articles acknowledges this connection. Citron & Pasquale, *supra* note 43, at 19; Crawford & Schultz, *supra* note 43, at 121–24.

226. Citron & Pasquale, *supra* note 43, at 19; Crawford & Schultz, *supra* note 43, at 122.

227. Citron & Pasquale, *supra* note 43, at 19.

228. *Id.* at 20.

229. Crawford & Schultz, *supra* note 43, at 124.

230. *Id.* at 126; Citron & Pasquale, *supra* note 43, at 26, 28. Companies could provide the information in a public statement that resembles a privacy policy. Crawford & Schultz, *supra* note 43, at 126. Or, they could make it available to affected individuals upon request. Citron & Pasquale, *supra* note 43, at 28.

231. Crawford & Schultz, *supra* note 43, at 127–28.

232. *Id.* at 126; *see also* Citron & Pasquale, *supra* note 43, at 20, 25, 31.

would review it for bias, errors or other problems.²³³ This would “protect [both the] scorers’ intellectual property and [the] individuals’ interests.”²³⁴

The purpose of these transparency requirements is to allow affected individuals to understand and, if necessary, challenge the algorithmic decisions.²³⁵ Each article argues that individuals should be given an opportunity for such a hearing. This could be a hearing before a representative of the company itself.²³⁶ Or, an individual could bring a claim before a neutral arbiter such as the FTC and allege that the determination was biased or had otherwise treated the person unfairly.²³⁷ The FTC (or other neutral arbiter) would investigate allegations of irregularities that “might render the adjudication unfair.”²³⁸

In addition to procedural protections, Professors Citron and Pasquale argue that companies should be required to obtain a license for their algorithms before using them to make impactful decisions about individuals.²³⁹ The licensor could either be a regulator such as the Equal Employment Opportunity Commission or a private entity that a regulator had itself licensed.²⁴⁰ Through the licensing process, public values would be brought to bear on these private business activities. Following the due process metaphor, regulatory oversight of corporate algorithmic adjudications to ensure substantive fairness²⁴¹ is akin to supplementing procedural due process with substantive due process.

Professors Citron and Pasquale also make another relevant proposal. They suggest that “[t]he FTC can oversee credit-scoring systems under its authority to combat ‘unfair’ trade practices under Section 5 of the [FTC] Act.”²⁴² Citron and Pasquale’s focus on transparency and other due process rights can cause the reader not to take proper notice of this point. But it is an important one. Here, Citron and Pasquale suggest that, with respect to credit scoring systems at least, the FTC could engage in *substantive* assessments as to whether the scoring models are fair. This insight, which Professor Mark

233. Citron & Pasquale, *supra* note 43, at 28; Crawford & Schultz, *supra* note 43, at 127.

234. Citron & Pasquale, *supra* note 43, at 28. While Citron and Pasquale focus their analysis on credit scoring systems, they explain that their approach “can extend more broadly to other predictive algorithms that have an unfair impact on consumers.” *Id.* at 20.

235. Crawford & Schultz, *supra* note 43, at 126–27.

236. *Id.* at 126–27. If trade secret or intellectual property concerns make the information too sensitive to share with such an individual, then a trusted third party could review it. *Id.*

237. Citron & Pasquale, *supra* note 43, at 28; Crawford & Schultz, *supra* note 43, at 127.

238. Crawford & Schultz, *supra* note 43, at 127.

239. Citron & Pasquale, *supra* note 43, at 21–22.

240. *Id.* at 22.

241. *Id.* Crawford and Schultz briefly mention “oversight and auditing primarily driven by public agencies” but do not further discuss or elaborate on this idea. Crawford & Schultz, *supra* note 43, at 124.

242. Citron & Pasquale, *supra* note 43, at 23; MacCarthy, *supra* note 44, at 425 (explaining that the FTC could use its unfairness authority to govern corporate data practices).

MacCarthy also articulated,²⁴³ sets the stage for this Article's more extended analysis of this approach.²⁴⁴

The technological due process approach to the governance of predictive analytics has some real strengths. By providing regulators and data subjects with access to the algorithms and data, this approach will bring sunlight into these eligibility determinations and so deter abusive, biased, and erroneous decisions. By giving individuals notice of and the opportunity to challenge algorithmic adjudications, it gives people a tool for protecting themselves in the algorithmic economy. It addresses some of the procedural unfairness harms associated with predictive analytics.²⁴⁵

But the technological due process approach requires the individual to petition for access to the algorithms and data and to challenge them before a company representative or external authority. Most people will lack the knowledge and/or resources to do so. Procedural due process rights could—like the notice and consent regime they seek to shore up—sound good in theory but fall short when individuals actually try to take advantage of them.

Professors Citron and Pasquale's pre-deployment licensing idea takes the burden off individuals and puts it on the company and the licensors. But doing so could create a rigid regulatory bottleneck for companies that must move quickly in order to compete. In an analogous area, the Clean Air Act's pre-construction permitting requirements for facilities that emit criteria air pollutants have created major delays for industrial operations that need to rebuild their facilities quickly in order to update their products and remain competitive.²⁴⁶ The algorithmic economy moves faster than the smokestack one, and such delays could take an even greater toll on competitiveness, especially if other countries did not require such licensing.

Finally, the licensing requirement and some of the articles' other proposals would require legislation which will be difficult to achieve in a gridlocked Congress.²⁴⁷

243. MacCarthy, *supra* note 44, at 488, 490–91 (explaining that the FTC could use its unfairness authority to govern privacy issues). In contrast to Professor MacCarthy, this Article would have unfairness authority govern all commercial uses of predictive analytics, not just those that qualify as “impermissible uses,” *id.* at 480, 482, and would not continue to rely heavily on “notice and affirmative consent,” *id.* at 496. Unlike Professors Citron and Pasquale, this Article does not advocate for pre-implementation licensing of algorithms, Citron & Pasquale, *supra*, note 43, at 22, and focuses more on substantive than on procedural solutions. That said, this Article owes a debt to Professor MacCarthy and to Professors Citron and Pasquale for their initial insight that the FTC could productively use its unfairness authority to govern privacy and data analytics. See MacCarthy, *supra* note 44, at 490–91; Citron & Pasquale, *supra* note 43, at 23.

244. See *infra* Part V.

245. See *supra* notes 121–123 and accompanying text.

246. Dennis D. Hirsch, *Lean and Green? Environmental Law and Policy and The Flexible Production Economy*, 79 IND. L. J. 611, 635–36 (2004).

247. Citron & Pasquale, *supra* note 43, at 22.

One of Citron and Pasquale's recommendations avoids these pitfalls: their suggestion that the FTC use its unfairness authority to evaluate algorithmic credit scoring systems.²⁴⁸ This puts the burden on regulators, who have the knowledge and experience to act, rather than on individuals, who often do not.²⁴⁹ If the punishments are significant enough, this oversight could deter unfair data practices going forward without producing a regulatory bottleneck. It makes sense to look to unfairness authority for a potential solution.

V. PREDICTIVE ANALYTICS AND THE FTC'S UNFAIRNESS AUTHORITY

The unfairness approach raises as many questions as it answers. Does the FTC have legal authority to apply its unfairness authority to analytics eligibility determinations, or would a statutory amendment be required after all? If the FTC does have this authority, how would unfairness authority apply to analytic adjudications? What does it mean for a business practice to be "unfair," and how would the FTC apply this test in this context? Would the FTC's use of its unfairness authority protect individual and broader social interests sufficiently? Would it give the five FTC Commissioners too much unfettered discretion to ban algorithmic practices they personally find objectionable? In a 2015 essay, the author began a more sustained analysis of unfairness authority and predictive analytics.²⁵⁰ This Part expands on this initial inquiry and explores, in far greater depth, whether the FTC could use its unfairness authority to regulate predictive analytics.²⁵¹ It begins by describing the FTC's unfairness authority and how it could apply to predictive analytics. To make the inquiry more concrete, it identifies a real-world data analytics example and explores how unfairness authority would apply to that situation. It then focuses on Congress's instruction to the FTC that, in making unfairness determinations, the FTC "may consider established public policies,"²⁵² and explains this phrase's importance. This Part concludes by proposing that the FTC use its unfairness authority to regulate predictive analytics.

248. *Id.* at 23.

249. Although a regulator such as the FTC would still confront the problem of limited resources with which to audit and enforce against algorithms.

250. Hirsch, *supra* note 45.

251. As explained above, *supra* notes 44 and 46, the author recognizes and builds on Professor Mark MacCarthy's, and Professors Danielle Keats Citron and Frank Pasquale's, important insight that the FTC could use its unfairness authority to govern data analytics.

252. Federal Trade Commission Act, 15 U.S.C. § 45(n) (2012).

A. *Unfairness Authority*

Section 5 of the FTC Act gives the FTC the power to declare unlawful those business acts or practices that are “unfair or deceptive.”²⁵³ As was explained above this creates two distinct authorities: deceptiveness authority, and unfairness authority.²⁵⁴ The FTC’s Section 5 “deceptiveness” authority is closely linked to the control paradigm. It holds businesses to the representations that they include in their notices and that data subjects rely on when they consent. The FTC’s Section 5 “unfairness” authority is different. It requires the FTC to assess whether particular business practices—such as corporate use of data analytics—are fair or not. It authorizes the regulator to draw substantive lines rather than just to enforce the representations that businesses have made to consumers.²⁵⁵

Given the U.S. legal system’s preference for individual over social control, one would expect that the FTC Act would privilege deceptiveness authority over unfairness authority and would permit the FTC to utilize unfairness authority only where individual control was not possible. That is, in fact, what the Act does. In Section 5, the Act specifies that the FTC may only utilize its unfairness authority to address those consumer injuries that are “not reasonably avoidable by consumers themselves.”²⁵⁶ In other words, if consumers can protect themselves by making marketplace choices, then the FTC should focus on supporting these decisions and limit itself to its deceptiveness authority. If, on the other hand, individuals cannot protect themselves—if the consumer injuries are “not reasonably avoidable by consumers themselves”²⁵⁷—then consumer choice, by definition, is not up to the task. Here, Section 5 empowers the FTC to substitute its own judgment for that of the defenseless consumer and to determine whether the business act or practice in question is, or is not, fair. Deceptiveness authority gives way to unfairness authority. As the FTC has explained:

Normally we expect the marketplace to be self-correcting, and we rely on consumer choice—the ability of individual consumers to make their own private purchasing decisions without regulatory intervention—to govern the market. . . . However, it has long been recognized that certain types of sales techniques may prevent consumers from effectively making their own decisions, and that corrective action may then become necessary. . . . [T]he Commission’s unfairness matters are brought under these circumstances. They are brought, not to second-guess the wisdom of particular

253. *Id.* § 45(a)(1).

254. *See supra* notes 47–49 and accompanying text.

255. For an informative and comprehensive description of the FTC’s unfairness authority, see Solove & Hartzog, *supra* note 75, at 638–43.

256. 15 U.S.C. § 45(n).

257. *Id.*

consumer decisions, but rather to halt some form of seller behavior that unreasonably creates or takes advantage of an obstacle to the free exercise of consumer decisionmaking.²⁵⁸

The FTC's deceptiveness authority and unfairness authority thus divide the world in two. Where individuals with accurate information can reasonably protect themselves through their market choices, deceptiveness authority governs and seeks to make sure they have accurate (non-deceptive) information with which to do so. Where, even with accurate information, individuals cannot reasonably protect themselves in the marketplace, unfairness authority takes over.²⁵⁹ It draws substantive lines between fair and unfair business behavior and so protects individuals where they cannot protect themselves.

This move from supporting market choices to creating substantive protections is precisely the shift that the rise of predictive analytics makes necessary. As was explained above, predictive analytics changes the digital economy from a marketplace in which individuals could, at least to some extent, make meaningful choices about whether to share their personal information, to one in which they clearly cannot.²⁶⁰ Today, a company that gets a consumer to provide their surface data can infer their latent, often far more sensitive, information. Since individuals cannot understand in advance what their data might reveal, they can no longer "reasonably avoid" any injury that such disclosure might cause.²⁶¹ In such a situation, the FTC should move from liberalist deceptiveness authority to more interventionist unfairness authority.

The FTC Act sets out the contours of the FTC's unfairness authority in Title 15 section 45(n) of the United States Code, entitled "Standard of proof; public policy considerations." Section 45(n) states as follows:

The Commission shall have no authority . . . to declare unlawful an act or practice on the grounds that such act or practice is unfair unless the act or practice 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. In determining whether an act or practice is unfair, the Commission may consider established public policies as evidence to be considered with all other evidence. Such

258. FTC, COMMISSION STATEMENT OF POLICY ON THE SCOPE OF THE CONSUMER UNFAIRNESS JURISDICTION (1980), *reprinted in* Int'l Harvester Co., 104 F.T.C. 949 (1984), at 1070, 1074 [hereinafter FTC, POLICY ON UNFAIRNESS].

259. Beales, *supra* note 52, at 5.

260. *See supra* notes 31–34 and accompanying text.

261. 15 U.S.C. § 45(n).

public policy considerations may not serve as a primary basis for such determination.²⁶²

This language creates a three-prong test. In order to exercise its unfairness authority the FTC must first demonstrate that: (1) the business act or practice in question causes “substantial injury to consumers”; (2) consumers themselves cannot “reasonably avoid[]” this injury; and (3) the consumer injury that the business practice creates is “not outweighed” by its “benefits to consumers or to competition.”²⁶³ In addition, section 45(n) makes clear that, in determining whether a given act or practice is unfair, the FTC “may consider established public policies” along “with all other evidence.”²⁶⁴

B. *Predictive Analytics and Unfairness Authority: An Example*

Algorithmic determinations of eligibility for jobs, credit, insurance, housing, school admissions, and other important goods and life opportunities can harm individuals significantly when tainted by bias or procedural unfairness. This makes algorithmic adjudications an important context in which to assess whether laws and policies can protect people in the algorithmic economy. Companies consider predictive algorithms to be trade secrets.²⁶⁵ Once in a while, however, the precise algorithm—the set of steps that the company employs to make the eligibility determination—does come to light. Such was the case when, in 2008, the FTC filed a complaint against CompuCredit Corporation and Jefferson Capital Systems (“CompuCredit”).²⁶⁶ CompuCredit²⁶⁷ was in the business of issuing “subprime” credit cards to individuals who could not qualify for standard credit cards.²⁶⁸ It employed a “behavioral scoring model” to predict which of its card holders were most likely to default on their credit card debt. It then reduced these individuals’ credit lines by half.²⁶⁹

262. *Id.*

263. *Id.*

264. *Id.*

265. Citron & Pasquale, *supra* note 43, at 5.

266. FTC v. CompuCredit Corp., No. 1:08-CV-1976-BBM-RGV, 2008 WL 8762850 (N.D. Ga. Oct. 8, 2008). The FTC followed its traditional approach and pursued CompuCredit on deceptiveness grounds. *Id.* But it could equally have brought the enforcement action under an unfairness theory.

267. The company has since changed its name to Atlanticus Holdings. See *CompuCredit Holdings Corporation to Change Name to Atlanticus Holdings Corporation*, YAHOO! FIN. (Nov. 20, 2012), <https://finance.yahoo.com/news/compucredit-holdings-corporation-change-name-220000104.html>.

268. Complaint for Permanent Injunction and Other Equitable Relief at 5, *CompuCredit Corp.*, No. 1:08-CV-1976-BBM-RGV, 2008 WL 8762850 [hereinafter *CompuCredit Complaint*].

269. *Id.* at 34.

CompuCredit probably used a scoring model similar in form to the one by which Target determined which of its customers were pregnant.²⁷⁰ It identified a target variable—high risk of credit card default—that it sought to predict among its existing card holders. It had a set of data that included this target variable—the purchasing and credit payment histories of its past customers. It found a pattern of credit card purchases that correlated to default on credit card debt. These correlating items, the proxies,²⁷¹ included the use of the card for cash advances or for payments to:

- direct marketing merchants
- marriage counselors
- personal counselors
- automobile tire retreading and repair shops
- bars and night clubs
- pool and billiard establishments
- pawnshops
- and massage parlors.²⁷²

Where CompuCredit saw that a customer had made purchases of this type, it inferred the person posed a high risk of default and cut that person's credit in half.²⁷³

Was this algorithmic business practice “unfair”? If so, was the entire behavioral scoring model unfair, or just the use of certain proxies? By revealing CompuCredit's scoring model, the FTC's CompuCredit enforcement action provides us with a concise, real-world example with which to explore how unfairness authority applies to a business use of predictive analytics-based eligibility determinations. The following discussion examines the three unfairness prongs—(1) substantial injury, (2) not reasonably avoidable by the consumer, and (3) not outweighed by countervailing benefits²⁷⁴—and examines how each would apply to the CompuCredit scoring model.

1. Substantial Injury

The first prong of the unfairness test asks whether the business act or practice has caused “substantial injury” to consumers.²⁷⁵ Such injuries can consist of monetary, economic, health-related, or other types of tangible

270. See *supra* note 1 and accompanying text. The FTC Complaint does not describe in detail how CompuCredit used predictive analytics.

271. MAYER-SCHÖNBERGER & CUKIER, *supra* note 7, at 53–55.

272. CompuCredit Complaint, *supra* note 268, at 34.

273. *Id.*

274. 15 U.S.C. § 45(n) (2012).

275. *Id.*

harm.²⁷⁶ Some argue that they may also comprise emotional distress and other intangible harms,²⁷⁷ although there is considerable dispute over this.²⁷⁸ The magnitude of the injury must be more than trivial.²⁷⁹ To assess this, the FTC may add together a number of discrete injuries to particular individuals and assess whether the aggregate injury is substantial.²⁸⁰ The injury cannot be speculative.²⁸¹ There must be a significant, though not a certain, risk of it occurring.²⁸² In sum, in order to constitute “substantial injury” the harm must be more than trivial, can be aggregated harm from many discrete individuals, and must not be speculative. Economic harms definitely count. Emotional ones may count.

The CompuCredit behavioral scoring model meets this test. It reduces by half the credit available to those it deems to be at high risk of credit default. This is a concrete, economic harm that would satisfy even the most restrictive definition of “injury.” The loss of half of one’s available credit is not “trivial” from the perspective of that individual. If one adds up these discrete individual injuries it becomes even clearer that they are substantial. Nor is the injury speculative. CompuCredit is virtually certain to conclude that some of its card holders present too high a risk of default, and to reduce their credit accordingly. That is the reason that it developed the behavioral scoring model. CompuCredit’s analytics-based credit eligibility model clearly creates a “substantial injury” and so meets the first prong of the section 45(n) unfairness test.

2. *Unavoidable by Consumers Themselves*

Under the second element of the section 45(n) unfairness test, consumers must not reasonably be able to avoid the injuries on their own.²⁸³ The premise here is that, where consumers are able to avoid injuries through their

276. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073; MacCarthy, *supra* note 44, at 425; Beales, *supra* note 52, at 5.

277. MacCarthy, *supra* note 44, at 484 (“Emotional distress, mental anguish, loss of dignity and other harms are not ruled out by this criterion, but they must be effects that all or most or reasonable persons would construe as genuine harms.”).

278. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073 & n.16 (noting that emotional harm may be sufficient only in an “extreme case”); Beales, *supra* note 52, at 5 (stating that emotional harms are generally insufficient).

279. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073 (stating that “[t]he Commission is not concerned with trivial or merely speculative harms”).

280. HOOFNAGLE, *supra* note 38, at 132; MacCarthy, *supra* note 44, at 61; Beales, *supra* note 52 at 5, 8–9 (recognizing this as enforcement against Internet “mousetrapping”); *see also* FTC v. Certified Merchant Servs., Ltd. 126 F. App’x 651 (5th Cir. 2005); FTC v. Zuccarini, No. CIV.A. 01–CV–4854, 2002 WL 1378421 (E.D. Pa. Apr. 9, 2002).

281. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073.

282. MacCarthy, *supra* note 44, at 60–61.

283. 15 U.S.C. § 45(n) (2012); HOOFNAGLE, *supra* note 38, at 132–33.

market choices, it would be paternalistic for the FTC to do so on their behalf.²⁸⁴ Regulatory action is appropriate only where there is an “obstacle to the free exercise of consumer decisionmaking.”²⁸⁵ This element seeks to separate those instances in which consumers can use market choices to protect themselves from those in which they cannot. The FTC’s use of its unfairness authority is appropriate only in the latter category situations.²⁸⁶

CompuCredit’s card holders could not have imagined that their purchases were revealing their credit default risk any more than the female shoppers at Target could have known that their purchases were revealing their pregnancy status. The problem lies in the fact that predictive analytics, by its very nature, infers latent information from surface data. When individuals provide the surface data, they cannot know what latent information they are also revealing.²⁸⁷ It follows that the card holders could not reasonably have avoided the injuries that they experienced. CompuCredit’s scoring model and algorithmic eligibility determinations more generally meet the second prong of the unfairness test.

3. Cost-Benefit Analysis

The third element of the section 45(n) unfairness test asks whether the activity’s harms are “outweighed by countervailing benefits to consumers or to competition.”²⁸⁸ The FTC and commentators have interpreted this to require a cost-benefit analysis.²⁸⁹ The FTC is to weigh the substantial injury

284. Beales, *supra* note 52, at 5–6 (discussing consumer’s prerogative to avoid or not avoid injury through market choices).

285. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1074; Solove & Hartzog, *supra* note 75, at 639; Beales, *supra* note 52, at 5–6 (explaining that the FTC does not substitute its paternalistic choices for those of the consumer, even where Commissioners may believe that consumer making a poor choice, e.g., fast food or fast cars).

286. MacCarthy, *supra* note 44, at 62–63 (explaining that unfairness authority protects people where they cannot protect themselves through their own actions, and that business use of data analytics can constitute such a situation); Beales, *supra* note 52, at 5 (addressing “seller behavior that unreasonably creates or takes advantage of an obstacle to the free exercise of consumer decision-making” (quoting FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073)).

287. MacCarthy, *supra* note 44, at 487 (“Using advanced analytical techniques [businesses] can often infer with a reasonable degree of probability whether the person had the characteristic. Reasonable efforts by individuals to protect themselves by withholding information might be useless in the context of ubiquitous data collection and powerful analytics.”).

288. 15 U.S.C. § 45(n); *see also* MacCarthy, *supra* note 44, at 484.

289. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073–74; HOOFNAGLE, *supra* note 38, at 133 (explaining how the FTC looks at whether harms are injurious in their net effects); MacCarthy, *supra* note 44, at 487 (“This is essentially a requirement to do an assessment of the benefits of an information practice as well as the costs.”); Beales, *supra* note 52 (explaining that the FTC has established the “cost/benefit analysis” test).

that the business practice imposes (the costs), against the benefits that it creates.²⁹⁰ In order to be deemed unfair, a business practice must be “injurious in its net effects.”²⁹¹ In making this calculation, the FTC considers not only the costs and benefits to the individual business and consumers before it, but also those to society as a whole.²⁹²

CompuCredit’s behavioral scoring model creates some meaningful benefits. Assuming the model to be accurate, it will enable the company to reduce the number and scale of defaults. This will make its business more profitable. It will also enable it to offer credit at a lower rate of interest and to a wider array of individuals. This will allow some “sub-prime” borrowers to experience the benefits of a credit card who would not otherwise have been able to do so.

The costs of CompuCredit’s actions, too, are significant. To begin with, those whose credit is cut in half suffer a financial injury. The costs to the broader society are even more substantial. Knowing that use of a credit card to purchase particular items caused a reduction in available credit²⁹³ could deter people from purchasing these items in the future. The scoring model could accordingly lead to fewer people patronizing tire retreading shops, marriage counselors, personal counselors, bars and nightclubs, and massage parlors. People who drive too long on worn tires can have accidents that hurt themselves and others. People who need marital or personal counseling but forego it because they believe that it will lead to a reduction in their available credit can end up divorced, or unhappy, or both. CompuCredit’s behavioral scoring model could, accordingly, impose significant costs on the individuals concerned, their children, and the broader society.

So, do the benefits of CompuCredit’s scoring model outweigh its costs? The answer is clearer for some proxies than for others. The costs of deterring people from patronizing bars and nightclubs, pool and billiard halls, pawnshops and massage parlors are likely to be small. The countervailing benefits outlined above likely outweigh them. But the call is much tougher when it comes to tire retreading and counseling services because deterring the purchase of these services imposes a real cost on individuals and the broader

290. Beales, *supra* note 52, at 4–5 (stating that the Section 5’s unfairness prong creates a net benefit test); David L. Belt, *Should the FTC’s Current Criteria for Determining “Unfair Acts or Practices” Be Applied to State “Little FTC Acts,”* ANTITRUST SOURCE, Feb. 2010, at 1, 3, 11.

291. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073; *see also* Solove & Hartzog, *supra* note 75, at 639.

292. *See* FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073 (explaining that, in evaluating the costs, the FTC considers “not only the costs to the parties directly before the agency, but also the burdens on society in general”); HOOFNAGLE, *supra* note 38, at 132 (explaining that the FTC “will consider costs to the business and consumer, burdens on society . . . , burdens on the flow of information, and incentives for innovation”); MacCarthy, *supra* note 44, at 487 (stating that “[t]he assessment of countervailing benefits has to be made at the social level”).

293. Card holders could learn this through an inquiry, or from the coincident timing of the reduction and a particular purchase.

society. How is the FTC to assess whether these costs outweigh the benefits? And is the answer the same for tire retreading services as for marital and personal counseling?

C. Established Public Policies

The FTC Act provides some guidance. As quoted above²⁹⁴ Congress stated that, in evaluating the three unfairness prongs, “the Commission may consider established public policies as evidence to be considered with all other evidence.”²⁹⁵ The history behind this phrase informs its meaning. It is worth taking a brief detour to explore this history before coming back to the question of whether CompuCredit’s scoring model is unfair.

Congress amended the FTC Act in 1938 to give the FTC the authority to declare unlawful any “unfair or deceptive act[] or practice[] in or affecting commerce.”²⁹⁶ But it was not until the 1964 Cigarette Rule Statement of Basis and Purpose (“the Cigarette Rule”)²⁹⁷ that the FTC first articulated the criteria that it would use to make such unfairness determinations. The FTC explained that, in determining whether a particular business practice was “unfair,” it would examine:

- (1) whether the practice, without necessarily having been previously considered unlawful, offends public policy as it has been established by statutes, the common law, or otherwise—*whether, in other words, it is within at least the penumbra* of some common-law, statutory, or other established concept of unfairness; (2) whether it is immoral, unethical, oppressive, or unscrupulous; (3) whether it causes substantial injury to consumers (or competitors or other businessmen).²⁹⁸

Applying this test, the FTC found that cigarette manufacturers, in failing to disclose the health effects of their product, had behaved unfairly.²⁹⁹

The FTC’s reference to “public policy” in the first part of the Cigarette Rule test suggests that the FTC will look to the underlying values—the “penumbras”—of existing statutes, regulations, and common law decisions in order to inform how it interprets and applies its own unfairness authority. The FTC is not saying that it will enforce these other laws³⁰⁰ since other

294. See *supra* note 262 and accompanying text.

295. 15 U.S.C. § 45(n) (2012).

296. Wheeler-Lea Act, Pub. L. No. 447, sec. 3, § 5, 52 Stat. 111–12 (1938) (codified as amended at 15 U.S.C. § 45(a)); see also Beales, *supra* note 52, at 2.

297. 29 Fed. Reg. 8324, 8354–55 (July 2, 1964); Belt, *supra* note 290, at 1; Beales, *supra* note 52, at 2.

298. 29 Fed. Reg. 8324, 8355 (emphasis added).

299. *Id.*

300. Solove & Hartzog, *supra* note 75, at 640 (clarifying that “[u]nfair conduct need not be a violation of any particular law”).

agencies and courts presumably do that. Rather, it is saying that it will consider the “public policies” at the heart of these other laws when interpreting and applying its own unfairness authority.³⁰¹ While the FTC’s interpretation of unfairness would change in later years, this idea of looking to the public policies and “penumbras” that underlie existing laws as a reference point for determining what was unfair, has remained an important theme.

In the years that followed its issuance of the Cigarette Rule, the FTC made sparing use of its unfairness authority.³⁰² That changed in 1972 with the Supreme Court’s decision in *FTC v. Sperry & Hutchinson Co.*³⁰³ in which the Court cited with approval the Cigarette Rule’s three-part test.³⁰⁴ The FTC took this as an invitation to hold business acts and practices to be unfair if they violated *any one* of the three articulated factors.³⁰⁵ In the 1970s the FTC, thus empowered, embarked on its most extensive use of unfairness authority to date.³⁰⁶

This spawned a backlash.³⁰⁷ The business community,³⁰⁸ scholarly commentators,³⁰⁹ and ultimately members of Congress³¹⁰ expressed concern over the fact that the FTC could deem a business practice to be unfair if it

301. *Id.* (describing how the FTC looked to the Telecommunications Act’s restrictions on the disclosure of individuals’ phone records in determining, for unfairness purposes, that people had an expectation of privacy in their phone records).

302. Beales, *supra* note 52, at 2.

303. 405 U.S. 233 (1972); *see also* Beales, *supra* note 52, at 2 (discussing how the FTC expanded its use of unfairness authority following the *Sperry* decision). The case concerned whether the FTC’s unfairness authority governed only antitrust-like violations or could reach beyond this to other oppressive business practices, here, an issuer of trading stamps decision to forbid individuals from exchanging their stamps with one another. *Sperry & Hutchinson Co. v. FTC*, 432 F.2d 146, 150–51 (5th Cir. 1970), *modified and remanded by* 405 U.S. 233 (1972). The Supreme Court upheld the broader interpretation, stating that the FTC, in applying its unfairness authority, acts “like a court of equity, [and] considers public values beyond simply those enshrined in the letter or encompassed in the spirit of the antitrust laws.” 405 U.S. at 244.

304. *Id.* at 244 & n.5.

305. Belt, *supra* note 290, at 2 (citing an example).

306. *Id.* (discussing Over-the-Counter Drug Rulemaking regarding the ban of advertising directed to children on grounds that such advertising was “immoral, unscrupulous and unethical”).

307. HOOFNAGLE, *supra* note 38 at 60–66 (describing the FTC’s KidVid rulemaking, the controversy it provoked, and how both Congress and the FTC reacted to this controversy); Belt, *supra* note 290, at 2; G. S. Hans, *Privacy Policies, Terms of Service, and FTC Enforcement: Broadening Unfairness Regulation for a New Era*, 19 MICH. TELECOMM. & TECH. L. REV. 163, 168 (2012) (discussing the “very public and contentious debate” over FTC’s unfairness authority that played out from the late 1960s through the 1980s).

308. Hans, *supra* note 307, at 168 (stating that, in the 1970s, “[t]he FTC’s unprecedented vigorous enforcement led to a backlash from businesses that did not respond well to increased government oversight and regulation”).

309. Belt, *supra* note 290, at 2 (citing David A. Rice, *Consumer Unfairness at the FTC: Misadventures in Law and Economics*, 52 GEO. WASH. L. REV. 1, 25 (1984); Teresa M. Schwartz, *Regulating Unfair Practices Under the FTC Act: The Need for a Legal Standard of Unfairness*, 11 AKRON L. REV. 1, 21 (1977)).

310. Beales, *supra* note 52, at 3.

either offended public policy (the “penumbras” of existing laws), *or* was immoral/unethical, *or* caused substantial injury to consumers. These critics were particularly disturbed by the fact that, in making such unfairness decisions, the FTC did not need to consider the practices’ benefits and weigh them against the harms.³¹¹ They argued that such a standard gave the FTC virtually unfettered discretion to declare any business practice unfair, and so unlawful,³¹² if a majority of five Commissioners agreed that it was “immoral” or offended their view of public policy. Members of Congress were so upset that at one point Congress withheld the FTC’s funding and forced it to shut down for a few days.³¹³

The FTC responded with a policy statement through which it sought to constrain its own unfairness authority to mollify its critics.³¹⁴ In its 1980 Policy Statement on Unfairness,³¹⁵ the FTC articulated, for the first time, the now-familiar three-part test for when a given business practice is unfair. It stated, “To justify a finding of unfairness the injury must satisfy three tests. It must be substantial; it must not be outweighed by any countervailing benefits to consumers or competition that the practice produces; and it must be an injury that consumers themselves could not reasonably have avoided.”³¹⁶

The 1980 Policy Statement also changed the role that public policy would play in unfairness determinations. The FTC no longer identified it as a separate basis for declaring a business act or practice to be unfair. Instead, the FTC explained that it would, for the most part, use public policy only to inform its application of the new, three-prong test.³¹⁷ That is, it would look to existing, established public policies, ascertain their underlying values (their “penumbras”), and then interpret “substantial injury,” “countervailing benefits,” and the other aspects of the three-part test in light of these values. The FTC further committed to relying only on “well-established” public policies.³¹⁸ Employed in this way, public policy would inform the FTC’s interpretation of the three-prong unfairness test, make sure that it was in line with

311. *Id.* at 1–3.

312. *Id.*

313. *Id.*; HOOFNAGLE, *supra* note 38, at 65–66 (explaining that the KidVid controversy of the late 1970s “threatened [the FTC’s] very existence”).

314. HOOFNAGLE, *supra* note 38, at 121; Hans, *supra* note 307, at 168.

315. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073.

316. *Id.*

317. *Id.*; Hans, *supra* note 307, at 169; Beales, *supra* note 52, at 3. The FTC in the 1980 Policy Statement did, however, leave open possibility that it could use public policy as an independent basis for unfairness determinations when “the policy is so clear that it will entirely determine the question of consumer injury.” FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1075; Beales, *supra* note 52, at 3. Congress foreclosed this possibility in the 1994 Amendments when it codified the three-prong test. Thereafter, the FTC was only to use public policy to inform its interpretation of the three unfairness criteria; it was not to employ it as an independent basis for action.

318. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1076.

the values that Congress and the courts had expressed, and so provide an “important check on the overall reasonableness of the Commission’s actions.”³¹⁹

In 1994, Congress enacted Title 15, section 45(n) of the United States Code, quoted above,³²⁰ which codified the now-familiar three-prong unfairness test.³²¹ The new section 45(n) also confirmed public policy’s new role. Henceforth, public policy could serve *only* as a basis for interpreting the three, specified prongs.³²² It could not provide independent grounds for finding a particular business practice to be unfair.³²³

This history makes it possible to understand more fully Congress’s section 45(n) instruction, referred to above,³²⁴ that “[i]n determining whether an act or practice is unfair, the Commission may consider established public policies as evidence to be considered with all other evidence. Such public policy considerations may not serve as a primary basis for such determination.”³²⁵ This language does not permit the FTC to use public policy as independent grounds for finding particular business practices to be unfair. But it does allow the FTC to consider established public policies and their underlying values as a reference point for interpreting the three prongs of the unfairness test. It thus turns public policy from an independent basis for unfairness determinations into an interpretative tool that tethers the FTC’s unfairness judgments to the values that Congress and the courts have articulated elsewhere.³²⁶ In his highly informative book on the FTC and its involvement in privacy regulation, Professor Chris Hoofnagle explains that “established public polic[ies],” used in this context, mean policies that are “widely followed, and embodied in statutes, judicial decisions, or the Constitution.”³²⁷

This takes us back to the interpretative question that CompuCredit’s behavioral scoring model presents. Is it “unfair” to cut someone’s credit in half because they patronize bars and nightclubs? Tire retreading and auto repair shops? Marital and personal counselors? Can “established public policies”

319. *Id.* at 1075 n.27.

320. *See supra* notes 262–264 and accompanying text.

321. Federal Trade Commission Act Amendments of 1994, Pub. L. No. 103-312, sec. 9, § 5, 108 Stat. 1691, 1695 (codified as amended at 15 U.S.C. § 45(n) (2012)); HOOFNAGLE, *supra* note 38, at 131.

322. 15 U.S.C. § 45(n) (“In determining whether an act or practice is unfair, the Commission may consider established public policies as evidence to be considered with all other evidence.”).

323. *Id.* (“Such public policy considerations may not serve as a primary basis for such determination.”).

324. *See supra* note 295 and accompanying text.

325. 15 U.S.C. § 45(n).

326. Solove & Hartzog, *supra* note 75, at 639 (explaining that the FTC uses established public policies to help it determine when consumer injury is substantial).

327. HOOFNAGLE, *supra* note 38, at 133.

tell us anything about how to assess the costs and benefits of these algorithmic eligibility determinations?

They can. As was explained above, one of the costs of CompuCredit's approach is that it might deter people from engaging in the proxy activities.³²⁸ While there may be reasons why it might be harmful to deter people from frequenting bars and nightclubs (fewer opportunities to socialize) or from purchasing tire retreading services (greater risk of a car accident), no established public policy specifically identifies the value of these services or the social cost of having people forego them.

The situation is quite different with respect to marital and personal counseling services. All fifty states and the District of Columbia have established a psychotherapist-patient privilege in order to encourage people who need counseling services to visit and speak openly to a counselor.³²⁹ The U.S. Supreme Court in *Jaffee v. Redmond*³³⁰ took note of this universal common law principle in deciding, under Rule 501 of the Rules of Evidence, to adopt the psychotherapist-patient privilege as a matter of federal law and to extend it to social workers in addition to psychologists and psychiatrists.³³¹ The Court explained that the psychotherapist-patient privilege serves the public interest by facilitating the appropriate treatment of those experiencing mental illnesses.³³² "The mental health of our citizenry, no less than its physical health, is a public good of transcendent importance."³³³

CompuCredit's denial of credit to those who purchase marital or counseling services, if widely adopted in the credit industry, could deter people from obtaining such counseling services, even when they badly need them. This would not only injure the individuals in question but, as the Supreme Court and all fifty states have recognized, it would hurt the public as a whole by damaging mental health which is a "public good of transcendent importance."³³⁴ Given the many public policies that encourage marriage and treat it as valuable, one can assume that marital counseling, no less than personal counseling, provides great value to society, and that an algorithmic formula that penalized people for seeking marital counseling would have a similarly large, negative social impact.

328. See *supra* note 299 and accompanying text.

329. See *Jaffee v. Redmond*, 518 U.S. 1, 13 (1996), (noting that "the existence of a consensus among the States indicates that 'reason and experience' support recognition of the privilege").

330. 518 U.S. 1 (1996).

331. *Id.*

332. *Id.* at 11.

333. *Id.*

334. *Id.*

“Established public policies,” as the FTC Act uses the term, includes not only federal constitutional and statutory law but also federal judicial precedents, state statutes, and state common law.³³⁵ The Supreme Court’s *Jaffee* decision and the common law’s universal embrace of the psychotherapist-patient privilege suggest that the public gets great value when people seek out counseling services and would suffer a great cost were they to be deterred from doing so.

The underlying value of these federal and state common law decisions—the “penumbra” of these laws, as the FTC called it in the 1964 Cigarette Rule³³⁶—informs the interpretative question that CompuCredit’s behavioral scoring model presents. It suggests that CompuCredit’s decision to cut the credit of those who purchase marital and personal counseling service, while it may produce some benefits, could cause major harm to individuals and the broader society. It is thus “injurious in its net effects,” and so unfair.³³⁷ By contrast, CompuCredit’s decision to reduce the credit of those who purchase bar and nightclub, massage parlor, or even tire retreading services are less costly to society and so may be outweighed by their corporate and social benefits. This would make CompuCredit’s use of the latter proxies “fair” but its decision to cut the credit of those who purchase marital or personal counseling “unfair” and so unlawful.

D. A Proposal

The CompuCredit behavioral scoring model thus provides an example of how the FTC could apply unfairness authority, structured by the three-prong analysis and grounded in the penumbras of established public policies, to algorithmic eligibility determinations. It also shows that the FTC’s authority is not as unbounded as it was in the 1970s when its use of this power proved so controversial. Since the 1994 amendments to the FTC Act, the FTC must show that the activity in question meets the three-pronged, statutory unfairness test.³³⁸ While the test itself—particularly the third prong’s weighing of costs and benefits—affords the FTC a lot of discretion, Congress has provided it with an interpretative guide for navigating its way through the grey areas. The Commissioners are not to rely on their own views of what would, or would not, be best for society. They are to look to “established public policies,” identify the values that inform and underlie these policies and apply *those* values to the situation before them.³³⁹

335. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1074–76 (referring to “statute[s], common law” and “judicial decisions” as sources of established public policy).

336. 29 Fed. Reg. 8324, 8355 (July 2, 1964).

337. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1073.

338. *See supra* notes 262–264 and accompanying text (describing this test).

339. Lawyers routinely engage in this type of reasoning when interpreting statutes or regulations. EDWARD H. LEVI, AN INTRODUCTION TO LEGAL REASONING (1949).

Of course, there will be close calls, instances when there are important public policies on both sides of the argument. Congress's attempts to constrain the FTC's unfairness discretion suggests that, in such instances, the FTC should allow the business activity to proceed until such time as it becomes clearer that the injuries outweigh the benefits.³⁴⁰ But, even with this caveat, unfairness authority could still provide important protections. It could prevent the use of predictive analytics where the costs to individuals and society clearly outweigh the activity's benefits. On these grounds, the FTC could deem to be unlawful those commercial uses of predictive analytics that invade privacy, result in bias against protected classes, are unduly manipulative, or are procedurally unfair. In finding these practices to be unfair, the FTC could refer to the penumbras of existing anti-discrimination, unfair business practice, or privacy laws—established public policies all.

Over time, the FTC's unfairness determinations could establish a framework for distinguishing responsible from irresponsible business uses of predictive analytics.³⁴¹ The FTC could create some much-needed rules of the

340. Beales, *supra* note 52, at 5 (stating that the FTC should not be in the business of second-guessing market outcomes where costs and benefits are closely balanced).

341. *Cf.* Solove & Hartzog, *supra* note 75, at 649–50 (explaining how, in other areas, the FTC begins by articulating general standards but gradually begins to develop more specific rules). The FTC's use of unfairness authority to regulate corporate data security practices, discussed above, *see supra* Part V, has generated an interesting debate over the merits of such an approach to administrative policymaking. On the one hand, Solove and Hartzog commend the FTC for a “common law” approach that begins with general data security standards and, through a series of enforcement actions, articulates over time more specific guidelines. Solove & Hartzog, *supra* note 75, at 657. On the other hand, Professor Gus Hurwitz maintains that such policymaking through enforcement is a sub-optimal way of producing agency policy and denies regulated parties fair notice and due process. Gus Hurwitz, *Data Security and the FTC's UnCommon Law*, 101 IOWA L. REV. 955, 997–1008 (2016).

As the author has argued previously, the law is clear: The FTC and other agencies have broad discretion to choose between rulemaking and adjudicative policymaking as a means of developing policy. Hirsch, *supra* note 45, at 360–61; *see also* SEC v. *Chenery Corp.*, 332 U.S. 194 (1947) [hereinafter “*Chenery I*”]. Professor Hurwitz admits as much (“*Chenery II* is still good law today.”), though he argues that the Supreme Court may soon defer less and require greater reliance on rulemaking. Hurwitz, *supra* at 962.

If indeed the Court is to narrow agencies' discretion in this important area, it should not do so with respect to the FTC's regulation of predictive analytics. In *Chenery II*, the Court articulated the reasons why agencies may need to proceed by adjudicative policymaking instead of rulemaking:

[T]he agency may not have had sufficient experience with a particular problem to warrant rigidifying its tentative judgment into a hard and fast rule. Or the problem may be so specialized and varying in nature as to be impossible of capture within the boundaries of a general rule. In those situations, the agency must retain power to deal with the problems on a case-to-case basis if the administrative process is to be effective And the choice made between proceeding by general rule or by individual, *ad hoc* litigation is one that lies primarily in the informed discretion of the administrative agency.

Chenery II, 332 U.S. at 202–03.

The Court's reasoning applies with striking clarity to the regulation of predictive analytics. The field of predictive analytics is evolving rapidly in many economic sectors at once. The FTC is not even close to understanding it well enough to promulgate workable, general rules. In such an area, an agency must use adjudicative policymaking's incremental, common law approach “if the

road for this growing and increasingly important part of the economy. Such rules would not only protect individuals in the algorithmic era. It would also make the commercial application of predictive analytics more socially acceptable and so more sustainable in the long term, and so allow society more fully to enjoy its many benefits.³⁴² For all of these reasons, this Article proposes that the FTC use its unfairness authority to establish substantive parameters as to which uses of predictive analytics are socially appropriate and fair, and which are not.

E. Comparing the Various Approaches

The unfairness approach protects individuals from the threats of predictive analytics more fully and more effectively than the other proposed regulatory frameworks described in Part IV. To begin with, it does not depend on Congress to pass a new statute the way that the information fiduciary or contextual integrity approaches do. Congress has already enacted the statute: Section 5 of the FTC Act. The FTC has the authority, right now, to start declaring abusive algorithmic determinations to be “unfair.”³⁴³ This is a major advantage at a time of a gridlocked Congress and unpredictable President.

Next, unfairness authority tasks a regulator—the FTC—with evaluating the fairness of algorithmic adjudications. It does not burden individuals with the need to articulate all of their preferences the way that Mundie’s use-based approach does.³⁴⁴ Mundie’s proposal suffers from the same shortcoming as the notice-and-consent-based control regime that it seeks to replace: It saddles individuals with an unrealistic burden that will prevent most of them from taking advantage of the protection that the proposal offers. The unfairness approach does not rely on individual action. It gives this task to a regulator, one that has a long record of enforcement and has developed substantial expertise with respect to the information economy and privacy.³⁴⁵

administrative process is to be effective.” *Id.* at 203. Over time, the FTC may acquire enough knowledge and expertise to produce rules that make sense and achieve their intended aims. As Professor Margot Kaminski has recently argued, it may even employ a sector-based, co-regulatory approach to this end. Margot Kaminski, *Binary Governance: Lessons from the GDPR’s Approach to Algorithmic Accountability*, 92 S. CAL. L. REV. 1529 (2019). But in this area, the FTC is wise to start with investigative reports, workshops, guidance, and adjudicative policymaking, not generally applicable and inflexible rules. Hirsch, *supra* note 45, at 360–61.

342. WORLD ECONOMIC FORUM, DIGITAL TRANSFORMATION INITIATIVE IN COLLABORATION WITH ACCENTURE (2018), <http://reports.weforum.org/digital-transformation/wp-content/blogs.dir/94/mp/files/pages/files/dti-executive-summary-20180510.pdf>.

343. For a discussion of whether the FTC’s Section 5 unfairness authority reaches this far, see *infra* Section VI.A; see also Hartzog, *supra* note 39, at 814 (“[T]he FTC does not need a new authorization of power to tackle a new technology. It is sufficient if a company uses a new technology in commerce to harm or mislead consumers.”).

344. See *supra* note 227 and accompanying text.

345. The FTC will face some significant challenges in playing this role. These include its cumbersome Magnuson-Moss rulemaking authority, limited resources, lack of civil penalty authority,

The unfairness approach applies far more broadly than Professor Balkin's information fiduciary standard.³⁴⁶ As was explained above, Balkin's proposal would apply only to those companies that have direct relationships with individuals and so can be said to stand in a fiduciary relationship with them.³⁴⁷ The Section 5 unfairness approach suffers from no such limitation and the FTC could apply it to all parts of the algorithmic economy. Balkin himself recognizes that his proposed approach does not reach those companies that make algorithmic eligibility determinations about individuals but have no relationship with them. For these companies, he acknowledges that:

[W]e can no longer rely on the notion of special fiduciary relationships between individuals and companies to regulate the use and abuse of data. Instead, we must ask what duties of good faith and ethical conduct in the collection, analysis, use, sale and distribution of data are owed to the members of society as a whole.³⁴⁸

That question—what duties of good faith and ethical conduct do all companies that employ predictive analytics owe to society—is precisely the one that unfairness authority wrestles with and seeks to answer.

The unfairness approach also avoids the other main problem with the information fiduciary model. As was explained above, a market system does not require a business (e.g., a company considering an employee for a promotion) to put individuals' interests *ahead* of its own.³⁴⁹ That is why the law has traditionally imposed fiduciary obligations only on a small subset of professions (e.g., lawyers, doctors) whose very role requires that they put their clients' interests first. It allows other companies to make appropriate business decisions so long as they do so in a way that is fair and does not unduly

and controversial history. See HOOFNAGLE, *supra* note 38, at 334–35 (describing these weaknesses). Congress could amend the FTC Act to remedy some of these shortcomings. Even without such legislation, it is worth remembering the strengths the FTC would bring to the task. As Professor Chris Hoofnagle articulated so well, “the FTC is a nimble agency” that is necessarily bi-partisan, operates with a common law approach that builds policies incrementally, combines economic considerations with other social priorities, and has shown itself to be a good strategist. *Id.* at 364–65. Similarly, Professor Woody Hartzog has explained that the FTC is well-suited to govern robots and other emerging technologies. Hartzog, *supra* note 39, at 824–31. The FTC:

[H]as developed a robust body of law to draw from . . . has a track-record of fostering nascent technologies like the Internet . . . gives deference to industry standards where relevant, which will keep the law . . . from being arbitrary and disconnected from practice . . . [and] regularly cooperates with other regulatory bodies and can use this experience to build consensus . . .

Hartzog, *supra* note 39, at 825. In looking for an agency to set the rules of the road for the algorithmic economy, one could do far worse.

346. See *supra* notes 176–181 and accompanying text.

347. Balkin, *supra* note 40, at 1233–34.

348. *Id.*

349. See *supra* notes 178–181 and accompanying text.

take advantage of others. That is what the FTC's unfairness authority requires. It makes far more sense to hold companies in the algorithmic economy to an unfairness standard than to duties of care or loyalty.

The unfairness approach would not produce a regulatory bottleneck of the type that Professors Citron and Pasquale's pre-deployment licensing requirement would create.³⁵⁰ Under its unfairness authority, the FTC would evaluate algorithmic determinations *after* a company makes them and would at that point determine whether or not they were unfair. It would not make companies wait for this determination before proceeding, as they would under the Citron and Pasquale proposal. The important role that speed and innovation play in the algorithmic economy makes *ex poste* unfairness authority much more practical than an *ex ante* licensing requirement. There is value in providing companies with *ex ante* guidance as to their obligations. But the FTC could achieve this without pre-deployment licensing. For example, the FTC could issue a report or other form of regulatory guidance to alert companies as to how it will apply the unfairness test to algorithmic decisionmaking. The growing bank of FTC unfairness determinations would further provide companies with a sense of their obligations without creating a regulatory bottleneck in this fast-moving and rapidly changing part of the economy.

The unfairness standard is much clearer and well-defined than Professor Cate, Professor Mayer-Schönberger's, and Peter Cullen's "cost-benefit" assessment,³⁵¹ or Professors Nissenbaum and Barocas's vague test for when norm-breaking algorithmic operations should be allowed.³⁵² Cate, Mayer-Schönberger, and Cullen do not explain how their cost-benefit analysis would work, calling instead on the legislature to pass laws that would "determine clearly how harms and benefits are to be evaluated."³⁵³ The FTC Act, by contrast, establishes the three-prong test and identifies a body of principles—"established public policies"—for interpreting these prongs.³⁵⁴ The FTC's decades of unfairness decisions further flesh out and define the unfairness standard.

Professors Nissenbaum and Barocas's invocation of broad principles like justice or autonomy would, in some respects, take us back to the days of the 1964 Cigarette Rule when the FTC found business acts to be unfair if, in the Commissioners' view, they were "immoral, unethical, oppressive, or unscrupulous."³⁵⁵ It was precisely such unbounded moral determinations that led to the backlash against the FTC in the 1970s.³⁵⁶ Since that time, the FTC

350. See *supra* note 246 and accompanying text.

351. See *supra* notes 207–208 and accompanying text.

352. See *supra* notes 195–196 and accompanying text.

353. CATE ET AL., *supra* note 10, at 17.

354. 15 U.S.C. § 45(n) (2012).

355. 29 Fed. Reg. 8324, 8355 (July 2, 1964).

356. See *supra* notes 307–313 and accompanying text.

and Congress have worked to define better the parameters of the FTC's unfairness authority through the 1980 Policy Statement on Unfairness and the 1994 FTC Act Amendments that created the three-prong test. Nissenbaum and Barocas's approach would cast off these constraints and, most likely, generate the type of controversy that the FTC faced in the past.

Finally, and perhaps most importantly, the FTC's unfairness authority possesses a political legitimacy that most of the other proposals lack. The new regulatory paradigm will need to draw substantive lines between predictive analytics practices that are socially acceptable, and those that are not. This will require engaging with, and choosing between, competing values. In a democracy, the elected representatives, or those to whom they have lawfully delegated their power, draw these lines most legitimately.³⁵⁷ Nissenbaum and Barocas's approach, which grounds its value judgments more in philosophical concepts than in legislative determinations, lacks political legitimacy in a democracy.

The FTC roots its unfairness determinations in the legislatively-determined three-prong test and in "established public policies." This positivist approach, which looks for values in existing legal texts, makes FTC unfairness determinations more politically legitimate than Nissenbaum and Barocas's philosophical judgments.

Professors Citron and Pasquale and Professors Crawford and Schultz focus, not on substantive value judgments, but on procedural mechanisms that will allow individuals and regulators to interrogate and challenge corporate algorithmic decisionmaking. This very useful addition should be built into any new paradigm for the regulation of predictive analytics. It assumes, rather than replaces, a standard for determining which algorithmic processes are socially acceptable, and which are not. It requires a means of drawing these substantive lines. As Citron and Pasquale themselves recognize, unfairness authority can provide this mechanism.³⁵⁸

VI. LEGAL AND POLICY QUESTIONS ABOUT THE UNFAIRNESS APPROACH

Part V identifies seven distinct advantages that unfairness authority has over the other main proposals.³⁵⁹ Specifically, unfairness authority: (1) does not depend on Congress to pass a new statute the way that the information fiduciary or contextual integrity approaches do;³⁶⁰ (2) tasks the FTC with making the fairness evaluations and does not require individuals to articulate all of their preferences the way that Mundie's use-based approach does;³⁶¹

357. Nemitz, *supra* note 197.

358. Citron & Pasquale, *supra* note 43, at 23.

359. *See supra* notes 343–358 and accompanying text.

360. *See supra* note 343 and accompanying text.

361. *See supra* notes 344–345 and accompanying text.

(3) applies far more broadly than Professor Balkin's information fiduciary standard;³⁶² (4) does not require businesses to put individuals' interests ahead of their own the way that the information fiduciary standard would;³⁶³ (5) provides guidance to companies without creating a regulatory bottleneck of the type that Citron and Pasquale's pre-deployment licensing would produce;³⁶⁴ (6) provides a clearer standard for distinguishing between acceptable and unacceptable analytic practices than Nissenbaum and Barocas's or Cate, Mayer-Schönberger, and Cullen's approaches would establish;³⁶⁵ and (7) proposes a more politically legitimate standard than Nissenbaum and Barocas's morality-based judgments.³⁶⁶ But an important question remains: does the FTC have the legal jurisdiction and political latitude to apply its unfairness authority to predictive analytics? This Part assesses whether the FTC's Section 5 unfairness authority is broad enough to encompass the regulation of predictive analytics, whether the FTC's use of its unfairness authority to regulate predictive analytics would offend the First Amendment, and whether such an application of unfairness authority would prove too controversial. It concludes that none of these issues should prevent the FTC from deploying its unfairness authority against unfair analytic practices.

A. Does the FTC's Section 5 Unfairness Authority Extend to Predictive Analytics?

Section 5 of the FTC Act authorizes the FTC to declare certain business acts and practices to be "unfair" and to order companies to stop engaging in them.³⁶⁷ Does this authority allow the FTC to pursue companies whose predictive analytics operations cause privacy invasions, manipulation, bias, and/or procedural unfairness? This is a question of statutory interpretation and should be settled by looking first to the statute's plain language, then to legislative history, and then to judicial interpretations of the statutory provision at issue.³⁶⁸ All three suggest that the FTC's unfairness authority is sufficiently broad and flexible to govern the risks that predictive analytics and other emerging technologies can create.

1. Plain Language

Beginning with the plain statutory language, FTC Act at Title 15, section 45(a)(1) of the United States Code states that "unfair or deceptive acts

362. See *supra* notes 346–348 and accompanying text.

363. See *supra* note 349 and accompanying text.

364. See *supra* note 350 and accompanying text.

365. See *supra* notes 351–354 and accompanying text.

366. See *supra* note 357 and accompanying text.

367. 15 U.S.C. § 45(a) (2012).

368. 73 AM. JUR. 2D *Statutes* § 83 (2012) (stating that legislative history aids in interpretation when a statute is not clear and unambiguous).

or practices in or affecting commerce, are hereby declared unlawful.”³⁶⁹ Section 45(n) does not affirmatively define the term “unfair” but rather specifies that the FTC may not deem a business act or practice to be unfair unless it meets the three-prong test described above.³⁷⁰ The Merriam-Webster Dictionary defines “unfair” as “1: marked by injustice, partiality, or deception : unjust; 2: not equitable in business dealings.”³⁷¹ Algorithmic decisions infected by bias are “marked by injustice”; the use of predictive analytics to manipulate people into transactions that harm them are “not equitable in business dealings;” and analytic processes that adjudicate people’s life opportunities without transparency, recourse, or other basic due process rights are “unjust.”³⁷² In short, algorithmic eligibility determinations can be “unfair” as the dictionary defines that term. As illustrated above in the CompuCredit example, they can also qualify as unfair under the FTC Act’s three-prong test.³⁷³ The statutory plain language appears to support the FTC’s use of unfairness authority to address predictive analytics’ potential harms.³⁷⁴

2. Legislative History

The term “unfair” first appears in the original, 1914 version of the FTC Act where Congress stated that “unfair methods of competition in commerce are hereby declared unlawful.”³⁷⁵ By the 1930s, questions began to arise about whether this authority served only an antitrust purpose and applied exclusively to businesses’ unfair actions against competitors or whether it also served a consumer protection purpose and governed companies’ unfair behavior towards consumers.³⁷⁶ The Wheeler-Lea Act of 1938 settled the question in favor of the latter view. It amended the FTC Act to read that “[u]nfair methods of competition in or affecting commerce, *and unfair or deceptive acts or practices in or affecting commerce*, are hereby declared unlawful.”³⁷⁷ This made clear that the FTC could enforce both against unfairly competitive practices and against unfair treatment of consumers. The word “unfair” thus appears twice in Section 5(a), once in reference to competition and the second

369. 15 U.S.C. § 45(a)(1).

370. See *supra* notes 262–264 and accompanying text.

371. *Unfair*, MERRIAM-WEBSTER, <https://www.merriam-webster.com/dictionary/unfair> (last visited Feb. 1, 2019).

372. See Citron & Pasquale, *supra* note 43, at 18–20.

373. See *supra* Section V.B.

374. Hartzog, *supra* note 39, at 803 (discussing the FTC’s longstanding use of its unfairness authority to regulate commercial manipulation of consumers).

375. FTC Act, ch. 311, 38 Stat. 717, 719 (1914) (codified as amended at 15 U.S.C. § 45(a)).

376. *LabMD, Inc. v. FTC*, 776 F.3d 1275 (11th Cir. 2015).

377. Wheeler-Lea Act, Pub. L. No. 447, sec. 3, § 5(a), 52 Stat. 111 (1938) (codified as amended at 15 U.S.C. § 45(a)(1)) (emphasis added).

time with regard to consumer protection. Courts have held that the word “unfair” has the same meaning in both places.³⁷⁸

Congress intended the FTC’s unfairness jurisdiction to be a “broad discretionary authority” that would enable the FTC “to define unfair practices on a flexible, incremental basis.”³⁷⁹ The reasons for this were straightforward. Congress recognized that unfairness could take many forms that the legislators themselves could not anticipate. As the Senate Report on the 1914 Act explained:

The committee gave careful consideration to the question as to whether it would attempt to define the many and variable unfair practices which prevail in commerce It concluded that . . . there were too many unfair practices to define, and after writing 20 of them into the law it would be quite possible to invent others.³⁸⁰

Given this, legislators recognized that the only way that the unfairness doctrine could retain its vitality over time would be for the FTC to have the flexibility to fit it to new business practices and harms as they arose.³⁸¹ “Congress affirmatively made a policy decision to choose vague language . . . because business practices and technology were constantly evolving, causing new problems that Congress could not quickly act to remedy.”³⁸² Congress believed that this flexibility would “endow the commission with tremendous discretion to move against abuses not yet invented.”³⁸³

3. Case law

Courts that have reviewed Section 5 and its legislative history have consistently reinforced the idea that Section 5 unfairness is broad, flexible, and

378. *FTC v. Wyndham Worldwide Corp.*, 10 F. Supp. 3d 602, 615 (D.N.J. 2014) (quoting *Am. Fin. Servs. Ass’n v. FTC*, 767 F.2d 957, 967 (D.C. Cir. 1985)); HOOFNAGLE, *supra* note 38, at 30.

379. *Wyndham Worldwide Corp.*, 10 F. Supp. 3d at 615 (quoting *Am. Fin. Servs. Ass’n*, 767 F.2d at 967). See generally HOOFNAGLE, *supra* note 38, at 10–13 (describing Congress’s decision to endow the FTC with flexible unfairness authority).

380. S. Rep. No. 63-597, at 13 (1914); see also STEPHANIE W. KANWIT, FEDERAL TRADE COMMISSION 113 (2012) (“[A]n enumeration, however comprehensive, of existing methods of unfair competition must necessarily soon prove incomplete, as with new conditions constantly arising novel unfair methods would be devised and developed.” (quoting *FTC v. Gratz*, 253 U.S. 421, 437 (1920) (Brandeis, J., dissenting))); accord HOOFNAGLE, *supra* note 38, at 10 (“The kinds of unfair behavior were too numerous to enumerate, and legislative prohibitions of them invited businesses to engage in practices that fell through minor loopholes.”).

381. KANWIT, *supra* note 380, at 113 (“The definition of ‘unfair’ in the original Federal Trade Commission Act and its amendments was deliberately left ambiguous and flexible.”).

382. HOOFNAGLE, *supra* note 38, at 119–20; accord KANWIT, *supra* note 380, at 71–72 (explaining that “[t]he voluminous legislative history . . . does indicate the following: (1) that the phrase ‘unfair methods of competition’ was deliberately left vague and open-ended, to depend on Federal Trade Commission interpretation” (footnote omitted)).

383. KANWIT, *supra* note 380, at 72; see also *LabMD, Inc. v. FTC*, 894 F.3d 1221, 1228 (11th Cir. 2018) (“Congress ‘intentionally left development of the term ‘unfair’ to the Commission’ through case-by-case litigation”); Hartzog, *supra* note 39, at 812–13.

capable of addressing new business practices and harms. They have stated that Congress designed the term as a “flexible concept with evolving content;”³⁸⁴ that it “intentionally left [its] development [of the term ‘unfair’] . . . to the Commission;”³⁸⁵ and that “Congress ‘explicitly considered, and rejected, the notion that it reduce the ambiguity of the phrase “unfair methods of competition” . . . by enumerating the particular practices to which it was intended to apply.’”³⁸⁶ In a 1934 case, the Supreme Court found it “unnecessary to attempt a comprehensive definition of the unfair methods which are banned, even if it were possible to do so. . . . New or different practices must be considered as they arise in the light of the circumstances in which they are employed.”³⁸⁷ These decisions established a pattern in which “courts have deferred to the Commission’s determination of what is unlawful in both the competition and consumer protection areas.”³⁸⁸

Recent FTC unfairness actions against companies that suffer data security breaches, and judicial decisions upholding this exercise of authority, reaffirm this long-standing judicial interpretation of the Act. A decade or so ago, the FTC began to focus on the growing number of data security breaches and their impact on consumers. It used its Section 5 unfairness authority to address this digital age threat. The FTC brought a series of complaints against companies that had suffered data security breaches on the grounds that, where consumers had entrusted a company with their personal information, a company’s failure to take reasonable measures to protect this data was “unfair.”

Most of the companies that received such a complaint settled with the FTC.³⁸⁹ Two, however, did not. In separate matters, Wyndham Worldwide Corp.³⁹⁰ and LabMD³⁹¹ challenged the FTC’s action in court. Each argued that the FTC had exceeded the bounds of its unfairness authority when it applied it to the corporate cybersecurity practices—an area of business action that did not even exist at the time of the 1938 Wheeler-Lea Act. Each of these cases thus squarely presented the question of whether the FTC’s unfair-

384. *FTC v. Bunte Bros., Inc.*, 312 U.S. 349, 353 (1941).

385. *FTC v. Wyndham Worldwide Corp.*, 799 F.3d 236, 243 (3d Cir. 2015) (first alteration in original) (quoting *Atl. Ref. Co. v. FTC*, 381 U.S. 357, 367 (1965)).

386. *Id.* (quoting *FTC v. Sperry Hutchinson, Co.*, 405 U.S. 233, 239–40 (1972)).

387. *FTC v. F.R. Keppel & Bro., Inc.*, 291 U.S. 304, 314 (1934).

388. HOOFNAGLE, *supra* note 38, at 30 (footnotes omitted).

389. GINA STEVENS, CONG. RESEARCH SERV., R43723, THE FEDERAL TRADE COMMISSION’S REGULATION OF DATA SECURITY UNDER ITS UNFAIR OR DECEPTIVE ACTS OR PRACTICES (UDAP) AUTHORITY 6–7 (2014) (stating that, since 2002, the FTC has settled twenty cases alleging that a company’s failure reasonably to protect consumer data constituted an unfair act or practice).

390. *Wyndham Worldwide Corp.*, 799 F.3d at 236.

391. *LabMD, Inc. v. FTC*, 894 F.3d 1221 (11th Cir. 2018).

ness authority was capacious enough to encompass digital age business practices and the threats that they caused. Each produced federal district court and circuit court decisions that addressed this question.

The Third Circuit in *FTC v. Wyndham Worldwide Corp.*³⁹² and the Eleventh Circuit in *LabMD, Inc. v. FTC*³⁹³ each held, consistent with the cases cited above,³⁹⁴ that Congress intended the FTC's unfairness authority to be broad and flexible so that the FTC could address new threats as they arose. After reviewing the history of the FTC's Section 5 unfairness authority the Third Circuit concluded that "[t]he takeaway is that Congress designed the term as a 'flexible concept with evolving content.'"³⁹⁵ It held that the FTC clearly had the power to declare unreasonably lax security measures that resulted in data security breaches to be unfair, so long as they satisfied section 45(n)'s three-prong test.³⁹⁶ The Eleventh Circuit, in *LabMD, Inc. v. FTC*, interpreted the statute in much the same way stating that, "[r]ather than list 'the particular practices to which [unfairness] was intended to apply,' Congress 'intentionally left development of the term "unfair" to the Commission' through case-by-case litigation."³⁹⁷ It, too, concluded that the FTC's unfairness authority was capacious enough to encompass unreasonably weak cybersecurity practices.

If the FTC's unfairness authority is broad enough to encompass unreasonably lax data security methods and the injuries they cause, it should similarly cover harmful predictive analytics practices and the threats they pose.³⁹⁸ Data security breaches and analytic unfairness each represents a new type of digital age threat that Congress could not have anticipated in 1938 when it passed the Wheeler-Lea Amendments but that it would have wanted the FTC to address through its flexible unfairness authority. If anything, predictive analytics poses the greater threat. Data security breaches injure privacy. But irresponsible use of predictive analytics can cause privacy invasions, manipulation, bias, and procedural unfairness. If the FTC's Section 5 unfairness authority is broad enough to include unreasonably lax data security measures, *a fortiori* it should encompass irresponsible predictive analytics practices. The Third and Eleventh Circuits' recent opinions on the FTC's unfairness authority are thus consistent with the FTC Act's plain language

392. 799 F.3d 236 (3d Cir. 2015).

393. 894 F.3d 1221 (11th Cir. 2018).

394. See cases discussed *supra* at notes 384–388 and accompanying text.

395. *Wyndham Worldwide Corp.*, 799 F.3d at 243 (quoting *Atl. Ref. Co. v. FTC*, 381 U.S. 357, 367 (1965)).

396. *Id.* at 244.

397. 894 F.3d at 1228 (alterations in original) (quoting *FTC v. Sperry & Hutchinson Co.*, 405 U.S. 233, 239–40 (1972)).

398. See *supra* Section V.B. (explaining that algorithmic business practices that invade privacy, produce biased results, manipulate privacy, and lack transparency and accountability mechanisms meet the three-prong unfairness test).

and legislative history and with earlier judicial determinations on this question.

In sum, the plain language, legislative history, and judicial interpretations of Section 5 each support a broad interpretation of the FTC's unfairness authority capable of meeting digital age threats.³⁹⁹ Were the FTC to declare harmful predictive analytics practices to be "unfair," a court would very likely uphold this exercise of its enforcement authority.

B. First Amendment Concerns

An FTC enforcement action against a company for its unfair use of predictive analytics should not offend the First Amendment. The key distinction here is between speech about matters of public concern and speech about private or market topics.⁴⁰⁰ The First Amendment offers far stronger protection for speech that shapes public opinion—public discourse—than for speech that is merely a form of personal or market behavior.⁴⁰¹ The latter category of speech acts "are not attempts to participate in the formation of public opinion by exchanging ideas, beliefs, and opinions. Instead, they are forms of market behavior that use speech. Therefore, states may regulate the speech involved in them."⁴⁰²

This Article is concerned primarily with predictive analytics-based eligibility determinations (algorithmic adjudications). Insofar as this data practice is speech at all,⁴⁰³ it is speech that is a form of market behavior. This kind of speech generally is not intended to shape public opinion.⁴⁰⁴ Consumer protection law can accordingly regulate it without running afoul of the First Amendment. As Professor Jack Balkin explained:

[T]he state can regulate what people say to each other as they form (or refuse to form) contracts, as it does in antitrust law, consumer protection law, and antidiscrimination law. . . . Therefore, without falling afoul of the First Amendment, governments can regulate contracts to prevent discrimination *and unfair business practices*; they can require companies to label their products and make disclosures to protect consumers; and they can require companies to disclose information about themselves and about their operations in order to protect investors.⁴⁰⁵

399. HOOFNAGLE, *supra* note 38, at 30 & n.78; Hartzog, *supra* note 39, at 813 (explaining that the "broad scope [of the FTC's unfairness authority] is ideal for a regulatory agency in charge of responding to challenges posed by new technologies").

400. Balkin, *supra* note 40, at 1211.

401. *Id.*

402. *Id.* at 1212.

403. *Id.* at 1194.

404. *Id.* at 1219.

405. *Id.* at 1213 (emphasis added).

Thus, it would appear that, much like fiduciary law,⁴⁰⁶ commercial unfairness law can, consistent with the First Amendment, regulate business collection and use of data for predictive analytics.

C. Is Unfairness Authority Too Controversial?

If the above analysis is correct, and the FTC can legally use its unfairness authority to address predictive analytics' threats, then why has it not already done so? The answer most likely lies in the FTC's past exercise of unfairness authority and the intense controversy that it provoked.⁴⁰⁷ As explained above,⁴⁰⁸ the FTC's expansive use of unfairness authority in the 1970s provoked howls of protest from industry representatives and members of Congress concerned that the five FTC Commissioners would use their unfairness authority to constrain any business behavior that they personally found to be inequitable or immoral. Congress even defunded the FTC for a few days just to make sure that it got the message.⁴⁰⁹ Would the FTC's use of unfairness authority to govern the rapidly expanding commercial use of predictive analytics provoke a similar reaction?

There are a number of reasons to believe that it would not. First, the legal landscape is different. The FTC limited the scope of its own unfairness authority when it issued the 1980 Policy Statement on Unfairness.⁴¹⁰ Congress further constrained the FTC when it amended the FTC Act in 1994 and added the three-prong unfairness test and the "established public policies" interpretative tool.⁴¹¹ Unlike the situation in the 1970s, a congressionally determined, statutory framework currently governs the FTC's exercise of its unfairness authority.

Second, as was explained above, the FTC has in the past decade successfully used unfairness authority to enforce against unreasonably lax corporate cybersecurity measures.⁴¹² Congress has acquiesced in this enterprising use of unfairness authority and the courts have approved it. The FTC could build on this precedent in moving to declare certain predictive analytics practices to be unfair and so unlawful.

406. *Id.* at 1216.

407. McSweeney, *supra* note 46, at 526 ("[T]he agency has, historically, run into significant resistance from industry and Congress when it is perceived as pushing the bounds of its authority to expand enforcement efforts innovatively. . . . The hangover from the so-called 'Kidvid' controversy remains a reminder to the FTC today that pushing too aggressively can result in painful consequences.").

408. *See supra* notes 307–313 and accompanying text

409. *See supra* note 313 and accompanying text.

410. FTC, POLICY ON UNFAIRNESS, *supra* note 258, at 1074.

411. *See* Federal Trade Commission Act Amendments of 1994, Pub. L. No. 103-312, sec. 9, § 5, 108 Stat. 1691, 1695 (codified as amended at 15 U.S.C. § 45(n) (2012)).

412. *See supra* notes 389–397 and accompanying text.

Finally, corporate behavior is already moving in this direction. Sophisticated companies that employ predictive analytics, machine learning, and AI are concluding that mere compliance with control-based privacy laws will not protect individuals and so will not protect the companies' own reputations. In recent years, industry-funded think tanks such as the Information Accountability Foundation and the Future of Privacy Forum have issued frameworks for the "ethical" practice of predictive analytics;⁴¹³ major companies have created a partnership committed to "fair and transparent" algorithms;⁴¹⁴ and "data ethics" has become a hot topic among corporate chief privacy officers.⁴¹⁵ Companies interested in such measures might not resist—indeed, some might even welcome—FTC guidance on how to draw the line between fair and unfair data practices. The FTC may well be able to use its unfairness authority to create rules of the road for predictive analytics without encountering the resistance that it experienced in the 1970s.

VII. CONCLUSION

In their classic 1890 article *The Right to Privacy*, Justice Louis Brandeis and Samuel Warren argued that the law protected a person's "right to be let alone."⁴¹⁶ They maintained that those who violated this right to privacy should be held liable in tort. This argument, it is worth noting, does not call on individuals to draw the line between acceptable and unacceptable collection and use of their data. It looks to legislatures and courts—the institutions that fashion tort law—to do so. Brandeis and Warren's theory of privacy and of privacy law is thus quite different from Professor Alan Westin's.⁴¹⁷ Westin defines privacy as individual control over personal information⁴¹⁸ and looks to the legal system to provide this control through notice, consent, and purpose limitations. By contrast, Brandeis and Warren define privacy as the "right to be let alone"⁴¹⁹ and look to public authorities to define its boundaries and protect them through tort law.⁴²⁰

413. *Big Data Ethics Initiative*, INFO. ACCOUNTABILITY FOUND. (2013), <http://informationaccountability.org/big-data-ethics-initiative/>; JULES POLONETSKY, OMER TENE & JOSEPH JEROME, FUTURE OF PRIVACY FORUM, BENEFIT-RISK ANALYSIS FOR BIG DATA PROJECTS (2014), https://fpf.org/wp-content/uploads/FPF_DataBenefitAnalysis_FINAL.pdf.

414. See PARTNERSHIP ON AI, <https://www.partnershiponai.org/> (last visited Nov. 25, 2019).

415. See Simon McDougall, *The Role of the Chief Privacy Officer in 2020*, CPO MAG. (Feb. 23, 2017), <https://www.cpomagazine.com/data-privacy/role-chief-privacy-officer-2020/>.

416. Warren & Brandeis, *supra* note 54, at 193.

417. Richards & Hartzog, *supra* note 11, at 436 (contrasting "tort" privacy with "control" privacy).

418. WESTIN, *supra* note 21, at 41–42.

419. Warren & Brandeis, *supra* note 54, at 195, 205.

420. *Id.* at 219; see also William Prosser, *Privacy*, 48 CALIF. L. REV. 383, 422 (1960) (defining four distinct privacy torts).

It is time to return privacy law to its roots. It is time to move away from an over-reliance on individual control and move to a complementary and equal focus on social protection, on public authorities drawing substantive lines between those data practices that are socially appropriate and fair and those that are unfair. Doing so will require hard value choices. Was Target's use of predictive analytics "unfair"?⁴²¹ Was CompuCredit's?⁴²² These are not easy issues to resolve. But if we want to live in a healthy way with the intensely powerful predictive technologies that increasingly order our lives, we have no choice but to ask these questions, wrestle with them, and democratically define the values that will structure the algorithmic economy.⁴²³ Individual control should continue to play a role. Where people can still make meaningful choices about the collection, use, and disclosure of their personal data,⁴²⁴ they should be given the opportunity to do so. But where they cannot, we do them no favors by asking them to make illusory choices. Where people cannot protect themselves, society needs to protect them. That is the lesson that landlord-tenant law learned so many years ago when it relinquished *caveat lessee* in favor of the implied warranty of habitability. It is the lesson that privacy law should learn today.

421. See *supra* note 1 and accompanying text.

422. See *supra* notes 266–273 and accompanying text.

423. Anita Allen, *Protecting One's Own Privacy in a Big Data Economy*, 130 HARV. L. REV. F. 71, 72 (2016) (calling for democratically-created limits on big data analytics).

424. Hartzog, *supra* note 12, at 954 ("The FIPS are necessary, but not sufficient.").