

2019

Power, Process, and Automated Decision-Making

Ari Ezra Waldman

Innovation Center for Law and Technology

Follow this and additional works at: <https://ir.lawnet.fordham.edu/flr>



Part of the [Law and Society Commons](#), and the [Science and Technology Law Commons](#)

Recommended Citation

Ari Ezra Waldman, *Power, Process, and Automated Decision-Making*, 88 Fordham L. Rev. 613 (2019).
Available at: <https://ir.lawnet.fordham.edu/flr/vol88/iss2/9>

This Symposium is brought to you for free and open access by FLASH: The Fordham Law Archive of Scholarship and History. It has been accepted for inclusion in Fordham Law Review by an authorized editor of FLASH: The Fordham Law Archive of Scholarship and History. For more information, please contact tmelnick@law.fordham.edu.

POWER, PROCESS, AND AUTOMATED DECISION-MAKING

Ari Ezra Waldman*

INTRODUCTION

Many decisions that used to be made by humans are now made by machines. Automated decision-making algorithms¹ evaluate teachers,² approve or reject loan applications,³ choose whom to search in an airport security line,⁴ allocate police officers on the beat,⁵ and determine eligibility for government benefits,⁶ among a litany of other commercial and government decisions.

To some, this is cause for celebration. Computers, the argument goes, are fast, powerful, and efficient and do not make the same mistakes as fallible,

* Microsoft Visiting Professor of Information Technology Policy, Princeton University, Center for Information Technology Policy; Professor of Law and Founding Director, Innovation Center for Law and Technology, New York Law School; Affiliate Fellow, Yale Law School Information Society Project. Special thanks to Denise Anthony, Danielle Keats Citron, Julie Cohen, Rebecca Crotoft, Deborah W. Denno, Ellen P. Goodman, Kristin Johnson, Margot Kaminski, Kirsten Martin, Frank Pasquale, Andrew Selbst, Olivier Sylvain, and Michael Veale. Thank you also to the outstanding members of the *Fordham Law Review* who helped organize an informative Symposium. This Essay was prepared for the Symposium entitled *Rise of the Machines: Artificial Intelligence, Robotics, and the Reprogramming of Law*, hosted by the *Fordham Law Review* and the Neuroscience and Law Center on February 15, 2019, at Fordham University School of Law. For an overview of the Symposium, see Deborah W. Denno & Ryan Surujnath, *Foreword: Rise of the Machines: Artificial Intelligence, Robotics, and the Reprogramming of Law*, 88 FORDHAM L. REV. 381 (2019).

1. In this Essay, I will use three phrases—“decision-making algorithms,” “automated decision-making,” “automated systems,” and derivations thereof—interchangeably to refer to the use of complex mathematical formulae to make commercial and social policy decisions.

2. See generally *Hous. Fed’n of Teachers, Local 2415 v. Hous. Indep. Sch. Dist.*, 251 F. Supp. 3d 1168 (S.D. Tex. 2017).

3. See Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 8–10 (2017) (describing the use of algorithms in credit scoring).

4. See Notice of Modified Privacy Act of 1974; Department of Homeland Security Transportation Security Administration—DHS/TSA—019 Secure Flight Records System of Records, 78 Fed. Reg. 55,270, 55,271 (Sept. 10, 2013) (“[T]he passenger prescreening computer system will conduct risk-based analysis of passenger data TSA will then review this information using intelligence-driven, risk-based analysis to determine whether individual passengers will receive expedited, standard, or enhanced screening . . .”).

5. See Andrew D. Selbst, *Disparate Impact in Big Data Policing*, 52 GA. L. REV. 109, 113–15 (2017) (describing predictive policing).

6. See generally MEREDITH WHITTAKER ET AL., AI NOW INST., AI NOW REPORT 2018, at 18–22 (2018), https://ainowinstitute.org/AI_Now_2018_Report.pdf [<https://perma.cc/6EFY-2FJP>].

arbitrary, ill-informed, and biased humans.⁷ In fact, the opposite is true. Automated decision-making systems based on “big data”-powered algorithms and machine learning are just as prone to mistakes, biases, and arbitrariness as their human counterparts.⁸ In some ways, these automated tools are worse: their opacity—even to experts, let alone ordinary citizens—makes algorithmic decisions difficult to challenge and analyze.⁹ And the faith we tend to put in the power of technology shields algorithmic systems from critical interrogation, in general.¹⁰ The result is a technologically driven decision-making process that seems to defy interrogation, analysis, and accountability and, therefore, undermines due process.¹¹ This should make algorithmic decision-making an illegitimate source of authority in a liberal democracy.¹²

7. See, e.g., Andrew Guthrie Ferguson, *Big Data and Predictive Reasonable Suspicion*, 163 U. PA. L. REV. 327, 389–95 (2015) (discussing some of the ways law enforcement can use large data sets to eliminate biases in policing); Ric Simmons, *Big Data, Machine Judges, and the Legitimacy of the Criminal Justice System*, 52 U.C. DAVIS L. REV. 1067, 1096–97 (2018) (noting that one of the “primary benefits of using predictive algorithms” is “their complete disregard of irrelevant subjective factors” like race, religion, what a person wears, how they conduct themselves in court, and so forth).

8. See, e.g., Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1256 (2007) (citing the example of Colorado’s Benefit Management System, which issued “hundreds of thousands of incorrect Medicaid, food stamp, and welfare eligibility determinations . . . since its launch in September 2004”).

9. See FRANK PASQUALE, *THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION* 1–17 (2015) (describing the ways in which data analyses about consumers are hidden from view and from legal process).

10. See, e.g., Ewart J. de Visser et al., *The World Is Not Enough: Trust in Cognitive Agents*, 56 HUM. FACTORS & ERGONOMICS SOC’Y ANN. MEETING 263, 266–67 (2012); P. de Vries & C. Midden, *Effect of Indirect Information on System Trust and Control Allocation*, 27 BEHAVIOUR & INFO. TECH. 17, 24, 27–28 (2008); Kevin Anthony Hoff & Masooda Bashir, *Trust in Automation: Integrating Empirical Evidence on Factors That Influence Trust*, 57 HUM. FACTORS 407, 413–15 (2015); Yung-Ming Li & Yung-Shao Yeh, *Increasing Trust in Mobile Commerce Through Design Aesthetics*, 26 COMPUTERS HUM. BEHAV. 673, 677–79 (2010); Richard Pak et al., *Decision Support Aids with Anthropomorphic Characteristics Influence Trust and Performance in Younger and Older Adults*, 55 ERGONOMICS 1059, 1070–71 (2012).

11. See Emily Berman, *A Government of Laws and Not of Machines*, 98 B.U. L. REV. 1277, 1279 (2018).

12. There are various definitions of “legitimacy” in the philosophical, political science, and legal literatures. For example, Mark Suchman defines legitimacy as “a generalized perception or assumption that the actions of an entity are desirable, proper, or appropriate within some socially constructed system of norms, values, beliefs, and definitions.” Mark C. Suchman, *Managing Legitimacy: Strategic and Institutional Approaches*, 20 ACAD. MGMT. REV. 571, 574 (1995). Tom Tyler, a law professor and political scientist, defines legitimacy as “perceived obligation to comply with the directives of an authority, irrespective of the personal gains.” TOM TYLER, *WHY PEOPLE OBEY THE LAW* 27, 45 (1990). A research agenda on the legitimacy of algorithmic decision-making must take a broad approach to the concept of legitimacy because the idea has traditionally been discussed in terms of authorities, individuals, or organizations. Algorithmic legitimacy can be based on the legitimacy of the authority, private or public, using it or on the legitimacy of the decision-making process or on the decision itself. Suchman’s broad definition, therefore, makes sense in this context. As such, I define legitimacy as the socially constructed propriety of authority to make decisions for others.

The question, then, is this: can systems that evade traditional modalities of accountability ever be legitimate? Perhaps they can't. Automating decisions about commercial and social goods may be at odds with important democratic values like equality, fairness, and human flourishing, whatever the benefits of computational efficiency may be. We see this approach, for example, in the rapidly spreading bans on the use of facial recognition technology.¹³

Or, as many scholars argue,¹⁴ perhaps we can leverage process and procedure to put guardrails around automated decision-making systems. This approach makes some sense. In his seminal study of legal legitimacy, Tom Tyler showed that popular perceptions of legitimacy and, in turn, a general willingness to accept the decisions of authorities hinge, at least in part, on the existence of procedural safeguards and the opportunity to be heard.¹⁵ In other words, legitimacy of legal authorities depends on process. This Essay challenges the application of this approach to automated decision-making.

Using algorithms to make commercial and social decisions is really a story about power, the people who have it, and how it affects the rest of us. Seen in this way, the project to limit individual and social harms caused by algorithms is different than scholars have so far presumed. In this Essay, I argue that algorithmic decision-making systems are social, political, and economic expressions of what Julie Cohen and others have called neoliberal managerialization, or an organizational system of public or private governance that prioritizes freedom and efficiency above all other values.¹⁶ Engineers, most of whom are heterosexual, white men,¹⁷ control the process by which decision-making policy is translated into decision-making code.¹⁸

13. See Kate Conger, Richard Fausset & Serge F. Kovalski, *San Francisco Bans Facial Recognition Technology*, N.Y. TIMES (May 14, 2019), <https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html> [<https://perma.cc/H2MP-ELVM>]; see also Woodrow Hartzog & Evan Selinger, *Facial Recognition Is the Perfect Tool for Oppression*, MEDIUM (Aug. 2, 2018), <https://medium.com/s/story/facial-recognition-is-the-perfect-tool-for-oppression-bc2a08f0fe66> [<https://perma.cc/S43S-DVZG>].

14. See *infra* Part II.A.

15. See TYLER, *supra* note 12, at 96, 116–20, 137–38, 149. The effect of fair procedure on legitimacy is more pronounced in contexts where procedure matters more, like a trial. *Id.* at 105.

16. See JULIE E. COHEN, BETWEEN TRUTH AND POWER: THE LEGAL CONSTRUCTIONS OF INFORMATIONAL CAPITALISM 7, 139–141, 143–46 (2019); see also LAUREN B. EDELMAN, WORKING LAW: COURTS, CORPORATIONS, AND SYMBOLIC CIVIL RIGHTS 124–50 (2016) (showing how form over substance in corporate compliance with civil rights law was having a deleterious effect on real progress on workplace equality).

17. See Kate Crawford, *Artificial Intelligence's White Guy Problem*, N.Y. TIMES (June 25, 2016), <https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html> [<https://perma.cc/J3G7-Z4WE>]; Karen Hao, *AI's White Guy Problem Isn't Going Away*, MIT TECH. REV. (Apr. 17, 2019), <https://www.technologyreview.com/s/613320/ais-white-guy-problem-isnt-going-away/> [<https://perma.cc/3XCQ-B32Q>].

18. See generally Ari Ezra Waldman, *Designing Without Privacy*, 55 HOUS. L. REV. 659 (2018) (describing how engineering control over the design process gives them power to determine the extent to which new technology designs protect privacy); see also WOODROW HARTZOG, PRIVACY'S BLUEPRINT: THE BATTLE TO CONTROL THE DESIGN OF NEW

And they exercise that power for corporate interests seeking greater efficiency under a deceptive discourse of neutrality and professional distance.¹⁹ In other words, algorithmic decision-making hides the fact that engineers and their corporate employers are choosing winners and losers while steadfastly remaining agnostic about the social, political, and economic consequences of their work. This is neoliberalism at its worst.

In this context, applying process-oriented solutions to algorithmic overreach won't work. The managerial ethos inside corporations operating in a permissive, neoliberal regulatory environment will twist process to serve corporate ends.²⁰ As such, regulators must go beyond process to rebalance the structures of power. Society should adopt normative or substantive mandates that all government and private automated decision-making must follow. Critically, evidence of achieving those normative goals cannot be limited to procedures put in place in their name, whether they be impact assessments, audits, or checklists.²¹ Evidence of algorithmic systems adhering to social values must come from independent interrogation of the model for noncompliance. We need to learn to audit the code, and legal rules must evolve to allow it. If these algorithms fail independent tests, they should be considered illegitimate in any liberal democratic society and not deployed.

This Essay proceeds as follows. Part I briefly describes and critiques the use of algorithmic systems to make decisions about us. This Part argues that the very elements that make algorithms so attractive to decision makers also make them illegitimate in a democratic society. Part II reviews existing proposals to rein in algorithmic decision-making, ultimately arguing that, in both theory and practice, neoliberal managerialization both creates the case for automated decision-making and undermines process-oriented forms of accountability. This Part concludes with a substantive approach to regulating algorithmic decision-making systems in society. The Essay concludes with avenues for future research.

I. ALGORITHMIC DECISION-MAKING AND ITS DISCONTENTS

Automated decision-making systems use complex mathematical algorithms to identify meaningful relationships and likely patterns in large data sets.²² For example, an algorithm can analyze diverse factors, including internet browsing behavior, purchase history, residence zip code, employment, educational achievement, salary, and family relationships,

TECHNOLOGIES 21–55 (2018) (recognizing the importance technology design plays in user autonomy or lack thereof).

19. See Waldman, *supra* note 18, at 681–96.

20. See EDELMAN, *supra* note 16, at 124–50 (discussing how compliance professionals inside corporations leveraged process and procedure to protect the company rather than achieve equality pursuant to Title VII of the Civil Rights Act of 1964).

21. See Ari Ezra Waldman, *Privacy Law's False Promise*, 97 WASH. U. L. REV. (forthcoming 2020) (manuscript at 3, 64), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3339372 [<https://perma.cc/SVN5-8R5B>] (showing that privacy protections can be undermined by process-oriented approaches to compliance).

22. See Berman, *supra* note 11, at 1279.

among myriad other variables, to predict the likelihood that someone with a particular collection of characteristics will be a productive employee, a good credit risk, an effective leader, or a qualified recipient of health care.²³ Humans may be able to roughly and imperfectly extrapolate leadership potential from a history of leadership on a resume or from an interview, but computers can both process more input variables and find unexpected relationships between inputs and dependent variables that humans could not.²⁴ That's what makes algorithmic decision-making so enticing to both public and private policymakers.

A growing literature explores some of the sociopolitical concerns with deploying automated systems to make decisions about our lives, particularly in the context of commerce, social policy, and criminal justice.²⁵ Indeed, the very characteristics that make automated decision-making systems so attractive—predictive abilities, complexity, power, and independence—are also what make them so problematic for the rule of law and legal legitimacy. Proposals aimed at making algorithms accountable to the law are attempts to address these problems. And yet, as I will show in the next section, the proposals are bandages that ignore the underlying incompatibility between algorithmic decision-making and a society based on normative values like equality and fairness.

A. Predictions and Probabilities

Algorithms cannot predict the future. They can, however, estimate the probability that something will happen based on existing data.²⁶ That is, algorithms cannot know for sure that a loan applicant will pay back her loan

23. For a more detailed discussion of how algorithms function on large data sets, please see Kate Crawford & Jason Schultz, *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*, 55 B.C. L. REV. 93, 96 (2014). See also FTC, BIG DATA: A TOOL FOR INCLUSION OR EXCLUSION? 1 (2016), <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf> [<https://perma.cc/8PTZ-ATHW>].

24. See David Lehr & Paul Ohm, *Playing With the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 666–71 (2017).

25. See, e.g., BRUCE SCHNEIER, DATA AND GOLIATH: THE HIDDEN BATTLES TO COLLECT YOUR DATA AND CONTROL YOUR WORLD 173 (2015) (noting that data is being collected to serve corporate and government ends); Steven M. Bellovin et al., *When Enough Is Enough: Location Tracking, Mosaic Theory, and Machine Learning*, 8 N.Y.U. J.L. & LIBERTY 555, 621–24 (2014) (arguing that law enforcement's use of location tracking could violate the Fourth Amendment); Kiel Brennan-Marquez, "Plausible Cause": *Explanatory Standards in the Age of Powerful Machines*, 70 VAND. L. REV. 1249, 1255–57 (2017) (challenging the use of algorithmic decision-making in the law enforcement context); Citron & Pasquale, *supra* note 3, at 4 (discussing how bias is embedded in algorithmic systems); see also Julia Angwin et al., *Machine Bias*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> [<https://perma.cc/4A9L-MLBW>] (showing racial bias in an algorithm used to predict recidivism).

26. See Bellovin et al., *supra* note 25, at 590 ("Machine learning works best when given a large training set of observations (ideally drawn in some independent manner) with which it estimates models. These models are then used to make predictions on future data outputting a probability measure for the occurrence of an event or existence of a fact.").

in time; they can only conclude that a certain mix of factors creates some probability that a result will happen. Like statistical analysis, then, automated decision-making systems will make mistakes. False negatives and misinterpretations will happen. Consider one popular example: researchers trained a machine to differentiate between images of dogs and wolves by feeding it images they had manually labeled “dog” and “wolf.” The program correctly classified many new images, but rather than learning about dogs and wolves, it found patterns that differentiated the pictures generally, particularly the presence of snow and trees. Wolves are far more likely than dogs to be found in the snow, so the algorithm identified all pictures with snow as “wolf.”²⁷

Sometimes, these mistakes can have real effects on the ground. In 2017, Amazon’s algorithm mistakenly recommended bomb-making products to be sold together.²⁸ Cancer victims have been erroneously denied benefits and have had to sue insurance companies for their due.²⁹ Welfare recipients have been stripped of entitlements.³⁰ Citizens have been placed on government watch lists for no reason other than an error in an automated system.³¹ Small mistakes can result in significant deprivation and harm, calling into question the extent to which algorithms can make fair, predictable decisions.

B. Complexity and Opacity

Mistakes are not the only threat to algorithmic legitimacy in a society that values the rule of law. As algorithms become more accurate and better predictors, they also become more complex and, therefore, more opaque and resistant to interrogation.³² Automated decision-making systems are “black boxes”;³³ even experts may not fully understand how inputs become outputs.³⁴ Algorithms are either intentionally kept secret, whether as

27. See Marco Tulio Ribeiro, Sameer Singh & Carlos Guestrin, “Why Should I Trust You?”: Explaining the Predictions of Any Classifier, 22 PROC. ACM SIGKDD INT’L CONF. ON KNOWLEDGE DISCOVERY & DATA MINING 1135, 1142 (2016).

28. See Paul Sandle, *Amazon Reviewing Website After Algorithm Suggests Bomb-Making Ingredients*, REUTERS (Sept. 20, 2017, 10:28 AM), <https://www.reuters.com/article/us-britain-security-amazon-com/amazon-reviewing-website-after-algorithm-suggests-bomb-making-ingredients-idUSKCN1BV1WK> [<https://perma.cc/K3TH-EAKY>].

29. See *generally* VIRGINIA EUBANKS, AUTOMATING INEQUALITY: HOW HIGH-TECH TOOLS PROFILE, POLICE, AND PUNISH THE POOR (2018).

30. See Scarlet Wilcock, *Policing Welfare: Risk, Gender and Criminality*, 5 INT’L J. FOR CRIME JUST. & SOC. DEM. 113, 114–15 (2016).

31. See Margaret Hu, *Big Data Blacklisting*, 67 FLA. L. REV. 1735, 1736 (2016).

32. See, e.g., Leo Breiman, *Statistical Modeling: The Two Cultures*, 16 STAT. SCI. 199, 206–08 (2001) (noting the trade-off between the accuracy and functionality of an algorithm, on the one hand, and the capacity of humans to understand how it works, on the other).

33. See PASQUALE, *supra* note 9, at 1–17.

34. See Edward K. Cheng, *Being Pragmatic About Forensic Linguistics*, 21 J.L. & POL’Y 541, 548 (2013).

proprietary trade secrets³⁵ or in order to prevent gaming,³⁶ or they are functionally opaque because of the “specialized knowledge” required to understand their source code.³⁷ Either way, the opacity of decision-making algorithms prevents those harmed by automated systems from determining either how a decision came about or the logic and reasoning behind it.³⁸ This makes accountability difficult.

C. Power and Privacy

Another advantage of automated decision-making is that algorithms can consider a multitude of factors—far more than humans³⁹—when analyzing enormous data sets and, therefore, find unexpected correlations between independent and dependent variables.⁴⁰ Algorithms powered by machine learning can also learn from experience and make even more accurate probabilistic determinations over time.⁴¹ This may be an algorithm’s chief advantage: it can do things with numbers that humans can’t.

But algorithms’ superhuman power to process and find unexpected insight in raw data also makes them privacy invasive and, arguably, inconsistent with democratic principles of autonomy, dignity, and choice. As the computer scientist Steve Bellovin has explained, “[m]achine learning algorithms are able to deduce information—including information that has no obvious linkage to the input data—that may otherwise have remained private due to the natural limitations of manual and human-driven investigation.”⁴² That can result in the revelation of intimate information, from membership in a protected class to an underage pregnancy.⁴³ It also undermines the obscurity

35. See Citron & Pasquale, *supra* note 3, at 5 (stating that algorithms are “shrouded in secrecy”). See generally Rebecca Wexler, *Life, Liberty, and Trade Secrets: Intellectual Property in the Criminal Justice System*, 70 STAN. L. REV. 1343 (2017) (describing how trade secrecy impacts the criminal justice system).

36. See Citron & Pasquale, *supra* note 3, at 26 (noting that policies for algorithmic transparency should consider the potential for gaming). This concern is exaggerated. See Ignacio N. Cofone & Katherine J. Strandburg, *Strategic Games and Algorithmic Secrecy*, MCGILL L.J. (forthcoming 2019) (manuscript at 38–39), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3440878 [<https://perma.cc/95B8-E944>] (arguing that the likelihood of gaming is dependent on social costs and, thus, far less likely than scholars have assumed).

37. See Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 FORDHAM L. REV. 1085, 1092–94 (2018) (arguing that understanding machine learning algorithms takes “specialized knowledge” and even with that knowledge, the basis of a decision is often still inscrutable).

38. *Id.* at 1099–101.

39. See, e.g., Graeme S. Halford et al., *How Many Variables Can Humans Process?*, 16 PSYCHOL. SCI. 70, 70, 75–76 (2005) (finding that human processing maxes out at four variables).

40. See Andrej Zwitter, *Big Data Ethics*, BIG DATA & SOC’Y, July–Dec. 2014, at 1, 3–4.

41. See Harry Surden, *Machine Learning and the Law*, 89 WASH. L. REV. 87, 89 (2014) (explaining how machine learning can make algorithms capable of adapting “to enhance their performance on some task through experience”).

42. Bellovin et al., *supra* note 25, at 558.

43. See Charles Duhigg, *How Companies Learn Your Secrets*, N.Y. TIMES (Feb. 16, 2012), <http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html> [<https://perma.cc/R4JU-GN3C>] (describing Target’s pregnancy prediction score).

we come to expect in modern society⁴⁴ and allows data collectors to develop eerily precise virtual portraits of consumers for the purposes of behavioral targeting and manipulation through “dark patterns.”⁴⁵ And none of these concerns can be alleviated by a promise to delete personally identifiable information (PII) because algorithmic systems can use proxies for PII or develop PII from innocuous raw data.⁴⁶

Therefore, automated decision-making represents a radical shift in the discourse of power. Language, Foucault argued, shapes our understanding and perceptions of legitimacy and legality.⁴⁷ Critical race theorists and feminist scholars have made similar arguments about the power of speech.⁴⁸ Similarly when the language we use to talk about and implement privacy changes, so too does the locus of power over privacy. Through the noise of a diverse privacy discourse,⁴⁹ the language and assumptions about privacy have always been accessible to consumers: we talk about privacy in terms of “anonymity,”⁵⁰ greater “control” over our information,⁵¹ barriers and

44. See HARTZOG, *supra* note 18, at 10–11; Woodrow Hartzog & Frederic Stutzman, *The Case for Online Obscurity*, 101 CALIF. L. REV. 1, 5–8 (2013).

45. “Dark patterns” are design tricks that manipulate user behavior in predetermined ways. See, e.g., Saul Greenberg et al., *Dark Patterns in Proxemic Interactions: A Critical Perspective*, 2014 PROC. CONF. ON DESIGNING INTERACTIVE SYSTEMS 523, 524; Arunesh Mathur et al., *Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites*, 3 PROC. ACM ON HUM.-COMPUTER INTERACTION (forthcoming Nov. 2019), <https://arxiv.org/pdf/1907.07032.pdf> [<https://perma.cc/5TTJ-W743>]. One example of a dark pattern is PayPal’s alleged use of various website design features to trick users into signing up for a PayPal credit card when they thought they were using their preexisting PayPal account. See Complaint at 7–8, *Consumer Fin. Prot. Bureau v. PayPal, Inc.*, No. 1:15-cv-01426 (D. Md. May 19, 2015), ECF No. 1.

46. See Paul Ohm, *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*, 57 UCLA L. REV. 1701, 1704 (2010) (noting that it is impossible for perfectly anonymous data to have any marketable uses).

47. See MICHEL FOUCAULT, *THE HISTORY OF SEXUALITY* 101 (Robert Hurley trans., Pantheon Books ed. 1978) (1976) (noting that “[d]iscourse transmits and produces power”); Gerald Turkel, *Michel Foucault: Law, Power, and Knowledge*, 17 J.L. & SOC’Y 170, 172 (1990) (describing Foucault’s argument on “discourses of domination”); see also Richard K. Sherwin, *Dialects and Dominance: A Study of Rhetorical Fields in the Law of Confessions*, 136 U. PA. L. REV. 729, 749–824 (1988) (using case law on confessions to note a change in power dynamics).

48. See, e.g., Charles R. Lawrence, *If He Hollers Let Him Go: Regulating Racist Speech on Campus*, 1990 DUKE L.J. 431, 444; see also MARGARET THORNTON, *DISSONANCE AND DISTRUST: WOMEN IN THE LEGAL PROFESSION* 79–105 (1996) (arguing that Foucault’s discourse of power is fundamentally a gendered dynamic); PATRICIA J. WILLIAMS, *THE ALCHEMY OF RACE AND RIGHTS* 61 (1991) (noting that the legacy of slavery is entrenched by “powerful and invisibly reinforcing structures of thought, language, and law”).

49. See ARI EZRA WALDMAN, *PRIVACY AS TRUST: INFORMATION PRIVACY FOR AN INFORMATION AGE* 13–45 (2018) (grouping seemingly conflicting visions of privacy into negative and positive conceptions); see also DANIEL J. SOLOVE, *UNDERSTANDING PRIVACY* 14–36 (2008) (reviewing some of the many different definitions of privacy).

50. This is particularly helpful for members of marginalized and stigmatized communities. See, e.g., Scott Skinner-Thompson, *Outing Privacy*, 110 NW. U. L. REV. 159, 162 (2015) (arguing that privacy should be understood as preventing intimate information from serving as the basis of discrimination).

51. See, e.g., JULIE C. INNESS, *PRIVACY, INTIMACY, AND ISOLATION* 56 (1992) (defining privacy as “control over a realm of intimacy”); ALAN F. WESTIN, *PRIVACY AND FREEDOM* 7 (1967) (defining privacy as the freedom to make individual disclosure choices); Steve

separation,⁵² solitude,⁵³ and trust.⁵⁴ Shifting that discourse into the language of technology—namely, code⁵⁵—empowers the technologists and technology companies that control means of policymaking on the ground. Where a society uses code to determine winners and losers, the discourse of law becomes the discourse of engineering. This disempowers consumers, who have no access to a technology-driven privacy discourse, and again delegitimizes automated decision-making in a democratic society.⁵⁶

D. Independence and Bias

Finally, automated decision-making systems not only have the capacity to do more than humans ever could, they also supposedly remove humans and their biases from decision-making processes. But the mathematical system analyzing the data is agnostic about the value of the underlying data. In other words, algorithmic decision-making systems do not ignore biased data, they end up cementing those biases in society. The racial,⁵⁷ gender,⁵⁸ and

Matthews, *Anonymity and the Social Self*, 47 AM. PHIL. Q. 351, 351 (2010) (defining privacy as making the choice to “control” and “manage” the boundary between ourselves and others).

52. See, e.g., CHRISTENA NIPPERT-ENG, ISLANDS OF PRIVACY 8 (2010) (arguing that privacy can be subjective and, for some, only attained when they go “off the grid”); Robert S. Lauder & Maxine Wolfe, *Privacy as a Concept and a Social Issue: A Multidimensional Developmental Theory*, 33 J. SOC. ISSUES 22, 23 (1977) (identifying “separation from the collective” as a source of privacy); Jeffrey Rosen, *Out of Context: The Purposes of Privacy*, 68 SOC. RES. 209, 217 (2001) (describing privacy as a “shield”); Edward Shils, *Privacy: Its Constitution and Vicissitudes*, 31 LAW & CONTEMP. PROBS. 281, 283 (1966) (describing “private life” as a “secluded life”).

53. See Samuel D. Warren & Louis D. Brandeis, *The Right to Privacy*, 4 HARV. L. REV. 193, 196 (1890).

54. See WALDMAN, *supra* note 49, at 51–52 (noting that trust allows us to share because it creates expectations of confidentiality and adherence to norms).

55. Some scholars note that the discourse of artificial intelligence is inherently hidden from us. See, e.g., Maayan Perel & Niva Elkin-Koren, *Black Box Tinkering: Beyond Disclosure in Algorithmic Enforcement*, 69 FLA. L. REV. 181, 186–90 (2017) (discussing the limits of transparency in algorithmic systems); see also Julie Brill, Comm’r, Fed. Trade Comm’n, Keynote Address Before the Coalition for Networked Information: Transparency, Trust, and Consumer Protection in a Complex World 8–9 (Dec. 15, 2015), https://www.ftc.gov/system/files/documents/public_statements/895843/151216cnikeynote.pdf [<https://perma.cc/K82R-9G4S>] (noting difficulties in making algorithms transparent and calling on companies to address fairness themselves).

56. See generally Waldman, *supra* note 21.

57. See, e.g., Solon Barocas & Andrew D. Selbst, *Big Data’s Disparate Impact*, 104 CALIF. L. REV. 671, 712 (2016) (discussing how zip codes can serve as a proxy for race, thus embedding race discrimination in an algorithmic system).

58. See, e.g., Amanda Levendowski, *How Copyright Law Can Fix Artificial Intelligence’s Implicit Bias Problem*, 93 WASH. L. REV. 579, 609 (2018) (discussing examples of gender bias in algorithmic systems); see also David Meyer, *Amazon Reportedly Killed an AI Recruitment System Because It Couldn’t Stop the Tool from Discriminating Against Women*, FORTUNE (Oct. 10, 2018), <http://fortune.com/2018/10/10/amazon-ai-recruitment-bias-women-sexist/> [<https://perma.cc/79J7-J7XP>]; Noam Scheiber, *Facebook Accused of Allowing Bias Against Women in Job Ads*, N.Y. TIMES (Sept. 18, 2018), <https://www.nytimes.com/2018/09/18/business/economy/facebook-job-ads.html> [<https://perma.cc/HUS2-S5YK>]; Tom Simonite, *Machines Taught by Photos Learn a Sexist View of Women*, WIRED (Aug. 21, 2017), <https://www.wired.com/story/machines-taught-by-photos-learn-a-sexist-view-of-women/> [<https://perma.cc/EXJ4-U34L>].

socioeconomic⁵⁹ biases of algorithmic systems are particularly rich areas of research and popular reporting. That important work need not be repeated here. Suffice it to say, algorithmic systems are only as good as the data on which they are based. An algorithm can be used to predict recidivism rates among criminals, but if the inputs are biased against persons of color, the algorithms will overestimate the recidivism risk of black people and underestimate the risks for white people.⁶⁰ Similarly, a predictive language algorithm can anticipate the probability that certain words will be used in tandem—like “Paris” and “France” will likely be paired as often as “Seoul” will be paired with “South Korea”—but associate “man” with “doctor” and “woman” with “homemaker” because the underlying data on which the algorithm was based reflected society’s biases about gender.⁶¹ Coupled with our tendency to trust the neutrality and accuracy of computers far more than that of humans,⁶² the biases of automated decision-making systems powered by artificial intelligence can entrench second-class citizenship for marginalized populations.

II. PROCESS AND SUBSTANCE: RESPONDING TO THE THREATS OF ALGORITHMIC DECISION-MAKING

Scholars and advocates have recognized the threat that automated decision-making systems pose to the rule of law. As such, a variety of proposals seek to fabricate a regulatory regime for seemingly unaccountable algorithms. But their breadth belies one overarching trend: they all use process and procedure to achieve their goal. In this Part, I describe some of those proposals and argue that because both algorithmic decision-making and faith in process emerge from the same neoliberal ethos, many of the proposals under discussion today will fail to address the underlying social and political dangers algorithms pose to society.

A. Process-Oriented Proposals

Danielle Keats Citron was among the first to call for replacing old forms of adjudication and rulemaking with reconceived systems that include audit trails, education for hearing officers on machine fallibility, detailed explanations, publicly accessible code, and systems testing, among other recommendations.⁶³ Andrew Selbst and Solon Barocas argue that a right to explanation of automated decisions entitles individuals to clarity about the

59. See CATHY O’NEIL, WEAPONS OF MATH DESTRUCTION: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY 7 (2016); Mary Madden et al., *Privacy, Poverty, and Big Data: A Matrix of Vulnerabilities for Poor Americans*, 95 WASH. U. L. REV. 53, 77, 97 (2017). See generally KHIARA M. BRIDGES, THE POVERTY OF PRIVACY RIGHTS (2017).

60. See, e.g., Angwin et al., *supra* note 25.

61. This example comes from Levendowski, *supra* note 58, at 580–81.

62. See, e.g., de Vries & Midden, *supra* note 10, at 18–19; Hoff & Bashir, *supra* note 10, at 420; Poornima Madhavan & Douglas A. Wiegmann, *Effects of Information Source, Pedigree, and Reliability on Operator Interaction with Decision Support Systems*, 49 HUM. FACTORS 773, 777 (2007).

63. See Citron, *supra* note 8, at 1305–13.

process behind a model's development.⁶⁴ Dillon Reisman, Jason Schultz, Kate Crawford, and Meredith Whittaker of the AI Now Institute recommend algorithmic impact assessments, modeled after environmental or privacy impact assessments,⁶⁵ to document and assess a system's fairness.⁶⁶ Tal Zarsky suggests that greater transparency would help because anyone adversely affected by an algorithmic decision maker has a dignitary right to "understand why."⁶⁷ Mary Madden, Michele Gilman, Karen Levy, and Alice Marwick call for procedural safeguards to prevent algorithmic-based discrimination against the poor.⁶⁸ Sonia Katyal recommends codes of conduct, impact assessments, and whistleblower protections to alleviate bias problems.⁶⁹ Joshua Kroll and a team of researchers argue that technological tools can weed out biases from within.⁷⁰ Margot Kaminski suggests that the General Data Protection Regulation⁷¹ (GDPR), Europe's comprehensive data privacy legislation, can offer robust protection from harms stemming from automated decision-making because it entitles data subjects to explanations about the "logic" behind any algorithmic system.⁷² A. Michael Froomkin and Meg Leta Jones, among other scholars, have called for keeping humans in the loop to act as a check on automation running amok.⁷³ And Lilian Edwards and Michael Veale argue that the "right to be forgotten," data

64. See Selbst & Barocas, *supra* note 37, at 1087.

65. See Selbst, *supra* note 5, at 175, 184, 188 (discussing impact assessments in the environmental context).

66. DILLON REISMAN ET AL., AI NOW INST., ALGORITHMIC IMPACT ASSESSMENTS: A PRACTICAL FRAMEWORK FOR PUBLIC AGENCY ACCOUNTABILITY 3 (2018), <https://ainowinstitute.org/aiareport2018.pdf> [<https://perma.cc/ZJ4Q-BHLS>] (stating that "[p]ublic agencies urgently need a practical framework to assess automated decision systems and to ensure public accountability").

67. See Tal Zarsky, *Transparency in Data Mining: From Theory to Practice*, in DISCRIMINATION AND PRIVACY IN THE INFORMATION SOCIETY 301, 317 (Bart Custers et al. eds., 2013).

68. See Madden et al., *supra* note 59, at 113–22.

69. See Sonia K. Katyal, *Private Accountability in the Age of Artificial Intelligence*, 66 UCLA L. REV. 54, 107–28 (2019).

70. See Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633, 662–72, 682–92 (2017).

71. See generally Regulation 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), 2016 O.J. (L 119) 1 (EU) [hereinafter GDPR].

72. See Margot E. Kaminski, *The Right to Explanation, Explained*, 34 BERKELEY TECH. L.J. 189, 199 (2019).

73. See, e.g., A. Michael Froomkin, Ian Kerr & Joelle Pineau, *When AIs Outperform Doctors: Confronting the Challenges of a Tort-Induced Over-Reliance on Machine Learning*, 61 ARIZ. L. REV. 33, 34 (2019) (arguing that health algorithms without humans in the loop would decrease care quality); Stephen E. Henderson, *A Few Criminal Justice Big Data Rules*, 15 OHIO ST. J. CRIM. L. 527, 533–34 (2018) (proposing that a human should always be the ultimate decision maker when algorithms are deployed in the criminal justice system); Meg Leta Jones, *The Right to a Human in the Loop: Political Constructions of Computer Automation and Personhood*, 47 SOC. STUD. SCI. 216, 217 (2017) (couching the "human in the loop" mandate as based on European respect for human dignity).

protection impact assessments, and certifications and privacy seals can help make algorithms more accountable.⁷⁴

All of these proposals have one thing in common: they use procedural guardrails to prophylactically shoehorn algorithmic decision-making into an accountability regime. The approach makes some sense; procedure has, after all, been called the “essence” of due process rights.⁷⁵ The opportunity to be heard by an impartial adjudicator is central to legitimate democratic authority.⁷⁶ And in order to hold someone responsible or liable for intentional or negligent harm or failures in design, common law accountability regimes need to know what happened, what steps were taken to avoid causing harm, and what was reasonable under the circumstances. Procedural requirements like algorithmic impact assessments, source code transparency, explanations of either the result or the logic behind it, and a human in the loop who can hear someone’s appeal move opaque automated systems closer to more familiar, and more accountable, decision-making regimes.

B. The Neoliberal Project of Algorithmic Decision-Making

But that will not solve the problem. Transparency, impact assessments, paper trails, and the traditional accountability mechanisms they support do not address the gaps in the underlying social and political system that not only lays the groundwork for algorithmic decision-making but sees its proliferation, despite its biases, errors, and harms, as a good thing. The central legitimacy problem of automated decision-making systems is not that they are more opaque, complex, or biased than humans. If that were the issue, accountability regimes that work for human decision makers would, with some modifications, work for algorithms.

Rather, algorithmic decision-making represents a distinctly neoliberal form of policymaking that is, both doctrinally and as a matter of practice, agnostic about its sociopolitical and economic implications. Its emphasis on efficiency tilts the scales toward machines over humans and undermines the effectiveness of procedural guardrails to ensure accountability. For any society that cares about values like fairness, nondiscrimination, and human rights, algorithmic decision-making is presumptively illegitimate until it can be shown to reflect more than just neoliberal values of innovation and efficiency. To do that, we must learn to audit the code and not simply be satisfied with compliance procedures.

74. See Lilian Edwards & Michael Veale, *Slave to the Algorithm?: Why a ‘Right to an Explanation’ Is Probably Not the Remedy You Are Looking For*, 16 DUKE L. & TECH. REV. 18, 67–80 (2017).

75. See Citron, *supra* note 8, at 1255.

76. See TYLER, *supra* note 12, at 96, 116–20, 137–38, 149.

1. The Neoliberal Crusade for Efficiency

Neoliberalism is a political philosophy that aims to replace and undermine a political system based on social justice and social welfare with a regime that is “characterized by private property rights, individual liberty, unencumbered markets, and free trade.”⁷⁷ Its proponents argue that neoliberalism maximizes personal well-being by fostering the ability of individuals to choose their version of the good life.⁷⁸ The practical manifestation of this philosophy is a political and economic system where organizations, both governmental and commercial, prioritize efficiency and market solutions to social problems. Neoliberal trade policy, for example, is characterized by low or no tariffs and the free movement of goods and services, likely without the friction of worker and environmental protections;⁷⁹ neoliberal free speech policy is deregulated and based on the “marketplace of ideas” metaphor rather than concern for social justice and marginalized populations.⁸⁰ In a similar vein, neoliberal organizational policy is managerial, or streamlined for the purposes of fostering efficiency, innovation, and profit over other values.⁸¹ Managerial policies include anything that would amplify efficiency in the means of production, whether that is eliminating the likelihood of expensive litigation, erasing cost centers, or outsourcing tasks to more efficient entities.⁸²

This is where algorithmic decision-making comes in and why it is a natural outgrowth of a neoliberal philosophy. A political system that wants to prioritize efficiency, agility, and nimble economic action over social justice is a natural ally of professionals, like computer scientists and computer engineers, whose traditional “cardinal virtue” is efficiency.⁸³ It makes sense, then, that neoliberal leaders would turn to engineers to make policymaking

77. David Harvey, *Neoliberalism as Creative Destruction*, 610 ANNALS AM. ACAD. POL. & SOC. SCI. 22, 22 (2007).

78. *Id.* at 22–23. In this respect, neoliberalism is based, in part, on Rawlian, neo-Kantian theory. See generally JOHN RAWLS, A THEORY OF JUSTICE (1971) (using Kantian rationality to develop a system of justice for a multicultural society).

79. See, e.g., John O. McGinnis & Mark L. Movsesian, *The World Trade Constitution*, 114 HARV. L. REV. 511, 521–27 (2000) (noting that it is “well established” that the dominant neoliberal free trade regime promotes social welfare).

80. See, e.g., Jedediah Purdy, *Neoliberal Constitutionalism: Lochnerism for a New Economy*, 77 LAW & CONTEMP. PROBS. 195, 198–203 (2014) (describing neoliberal interpretations of constitutional provisions); Amanda Shanor, *The New Lochner*, 2016 WIS. L. REV. 133, 145 (chronicling a corporate social movement to interpret the First Amendment in line with neoliberal, libertarian ideas).

81. See Gerard Hanlon, *The First Neo-Liberal Science: Management and Neo-Liberalism*, 52 SOCIOLOGY 298, 309–11 (2016); see also COHEN, *supra* note 16, at 145 (noting that “[n]eoliberal managerialism does not value bureaucracy but rather efficient administration of lean and nimble production”).

82. See COHEN, *supra* note 16, at 144–57 (discussing outsourcing, forced arbitration, and the elimination of opportunities for plaintiffs to obtain justice in the courts as part of the managerial project).

83. Paul Ohm & Jonathan Frankle, *Desirable Inefficiency*, 70 FLA. L. REV. 777, 786 (2016). Ohm and Frankle also note that software engineers and computer scientists are increasingly being asked to consider other values in their work. See *id.* at 778–81.

more efficient.⁸⁴ Similarly, corporations influenced by neoliberal managerialism outsource to technology any activities that their human employees either do poorly or inefficiently.⁸⁵ For both reasons, then, automated decision-making systems seem attractive. Asking an algorithm to decide police allocation, jail time, or benefit entitlements reflects neoliberalism's underlying philosophical and political orientation toward deregulation, nonintervention, and efficiency. For profit-maximizing corporations looking to cut costs, replacing humans with mechanical decision makers makes sense from financial and efficiency perspectives as well.

2. Neoliberalism and the Power of Engineers

Algorithmic decision-making is, therefore, a natural outgrowth of neoliberal managerialism in the public and private sectors. This creates systemic problems for algorithmic legitimacy. Most notably, important social values like fairness, equality, nondiscrimination, accountability, fundamental human rights, and the extent to which society cares about the vulnerable and marginalized are absent from this narrative. This is the result of both doctrine and practice.

Neoliberalism assesses success or failure on the metrics and values of the market and, therefore, elides other normative social values by recasting or redefining them. Consider equality. Those who share a robust concept of equality recognize the obvious gap between formal law and life on the ground—where a law that, say, mandates nondiscrimination on the basis of race or gender fails to address the systemic burdens faced by marginalized populations and the hurdles they must jump over to achieve the same level of success as straight, white men. Neoliberalism recasts the project of equality to erase this problem. For neoliberals, equality is about equal opportunity or choice, and any gap between neoliberal equality and social practice is not the fault of the system but the fault of the individual. In the context of algorithmic decision-making, neoliberalism conceptualizes equality in terms of equal opportunity,⁸⁶ which is something machines do quite well. Computers neither know nor care if someone is black or brown or gay or transgender. Everyone is a number or a series of inputs, and code does not discriminate. This, of course, is a myth. As noted earlier, automated decision-making systems embed the biases of their inputs into their probabilistic determinations, cementing implicit and explicit prejudice in society.⁸⁷ As a result, true social equality is lost.

84. See COHEN, *supra* note 16, at 140 (bringing technology into government reflects the managerial ethos). We see this reflected in society and politics as well, where candidates for office highlight their corporate experience as a virtue, vow to cut “red tape,” and, at least in one instance, call themselves “the first MBA president.” See James P. Pfiffner, *The First MBA President: George W. Bush as a Public Administrator*, 67 PUB. ADMIN. REV. 6, 7 (2007).

85. See COHEN, *supra* note 16, at 156 (outsourcing to technology represents a focus on efficient management).

86. See Nancy Fraser, *Feminism, Capitalism and the Cunning of History*, 56 NEW LEFT REV. 97, 99, 108–11 (2009) (critiquing the feminist movement's slide toward neoliberalism).

87. See *supra* Part I.D.

Moreover, as a matter of social practice, algorithmic decision-making empowers engineers to make policy decisions, embedding their ingrained commitment to efficiency and their indifference to privacy and other social values in society. Engineers are the ones responsible for shoehorning policy into codable algorithms,⁸⁸ a translation process that will necessarily edit out the flexibility and contextualism that gives decision-making standards advantages over bright-line rules.⁸⁹ As such, a design team becomes the locus at which law is interpreted, negotiated, and transformed into actionable tools.⁹⁰ This gives engineers enormous power, both to choose winners and losers and decide what the law will mean in practice.⁹¹ And engineers leverage this power while professing that they and their designs are value-neutral. The law and technology scholar Frank Pasquale has noted that, even when programmers create tools that have evident racial biases—like a search engine that shows racially stereotyped advertisements⁹²—the designers cast the technology as a “cultural voting machine, merely registering, rather than creating, perceptions.”⁹³ Kate Crawford and Jason Schultz have described engineers as inadequately fluent in the language of public policy; as one programmer told them, “we can make stuff work—it’s not our job to figure out if it’s right or not. We often don’t know.”⁹⁴ And elsewhere, I have described how some engineers, at least in the high technology sector, resist the notion that designing for privacy or safety is part of their job or even possible under the demanding circumstances in which they work.⁹⁵ Therefore, when policy decisions are made by engineers, nonengineering values may get short shrift.

88. See Berman, *supra* note 11, at 1329–30.

89. Discussions about the relative merits of rules versus flexible standards are beyond the scope of this Essay. Duncan Kennedy originally described rules and standards as setting up a dialectical form of argument. See Duncan Kennedy, *Form and Substance in Private Law Adjudication*, 89 HARV. L. REV. 1685, 1689–90 (1976). Ronald Dworkin emphasized the role that standards play in realizing substantive legal principles. See Ronald M. Dworkin, *The Model of Rules*, 35 U. CHI. L. REV. 14, 22–29 (1967) (distinguishing between principles and rules in order to explain the important role of standards that are not rules); see also, e.g., MARK KELMAN, A GUIDE TO CRITICAL LEGAL STUDIES 15–63 (1987); RICHARD POSNER, THE PROBLEMS OF JURISPRUDENCE 42–53 (1990); FREDERICK SCHAUER, PLAYING BY THE RULES: A PHILOSOPHICAL EXAMINATION OF RULE-BASED DECISION-MAKING IN LAW AND IN LIFE (1991); Louis Kaplow, *Rules Versus Standards: An Economic Analysis*, 42 DUKE L.J. 557, 596–99 (1992); Pierre Schlag, *Rules and Standards*, 33 UCLA L. REV. 379, 383–420 (1985) (examining the form and rhetoric of the rules versus standards debate).

90. See generally Waldman, *supra* note 18.

91. See generally Ari Ezra Waldman, *Outsourcing Privacy* (May 20, 2019) (unpublished manuscript) (on file with author) (showing that engineers at privacy technology vendors code their interpretations of legal privacy requirements into technological tools).

92. See Latanya Sweeney, *Discrimination in Online Ad Delivery*, COMMS. ACM, May 2013, at 44, 46–47, 50–51.

93. PASQUALE, *supra* note 9, at 39.

94. Crawford & Schultz, *supra* note 23, at 105.

95. See Waldman, *supra* note 18, at 686–89.

3. The Failure of Process

Neoliberalism also explains why process-oriented accountability proposals will fail. In the neoliberal system that created algorithmic decision-making, accountability is recast as compliance, which can undermine the power of procedure to rein in automated decision-making in two ways.

First, procedure can be co-opted by corporate interests. According to the sociologist and legal scholar Lauren Edelman, the managerial law inside a corporation involves shifting the site at which law is interpreted and negotiated from policymakers and courts to corporate compliance procedures and internal structures designed to signify superficial legal adherence to regulations.⁹⁶ When that happens, procedural requirements can be co-opted to serve corporate, rather than consumer, interests. Edelman calls this the “mobilization of symbolic structures.”⁹⁷

In her book, *Working Law*, Edelman describes how some legal regimes characterized by vague requirements and process-oriented safe harbors give compliance professionals on the ground the opportunity to frame the law in accordance with managerial values like operational efficiency and reducing corporate risk rather than the substantive goals the law is meant to achieve, like consumer protection or equality.⁹⁸ This opens the door for companies to create structures, policies, and protocols that comply with the law in name only.⁹⁹ As these symbolic structures become more common, judges and policymakers defer to them as paradigms of best practices, mistaking mere symbols of compliance with adherence with legal mandates.¹⁰⁰ When this happens, law fails to achieve substantive goals because the compliance metric—the adoption of symbols, processes, procedures, and policies within a corporate environment—can be orthogonal to actual progress. Edelman discussed legal endogeneity in the context of race and sex discrimination in the workplace, where the equality goals of Title VII of the Civil Rights Act of 1964 were being frustrated by the ineffectual trainings, toothless policies, checklists, and disempowered diversity offices that compliance professionals created on the ground.¹⁰¹ Elsewhere, I have shown how this process is undermining the promised protections of data privacy laws, as well.¹⁰²

Even rules that mandate that all automated decision-making systems treat all individuals equally and fairly, but use process and procedure to do it, are likely to fall victim to the same phenomenon. Algorithmic impact assessments can identify and evaluate risks, consider alternatives, identify strategies to mitigate risks, and help articulate the rationale for the automated

96. See EDELMAN, *supra* note 16, at 100–50.

97. *Id.* at 153–55.

98. *Id.* at 3–15.

99. *Id.* at 14.

100. *Id.* at 12–13.

101. *Id.* at 11.

102. See generally Waldman, *supra* note 21.

system,¹⁰³ but they can also be mobilized, as Edelman argued in the nondiscrimination context, as a paper trail to push back against claims of unfair harm from those adversely affected by the algorithm.¹⁰⁴ Transparency, whether in the form of source code publication or an explanation of the results, can throw some sunshine on an opaque process but is functionally unhelpful to most individuals without specialized knowledge or convenient evidence for a fact finder to determine compliance with the law.¹⁰⁵ And keeping humans in the loop of a decision-making process offers an override to a mechanical process, but the extent to which the human element will have power depends entirely on the way the safeguard is implemented on the ground.

Second, the focus on documentation and process as ends in themselves elevates a merely symbolic structure to evidence of actual compliance with the law, obscuring the underlying substantive values of fairness, equality, and human dignity eroded by large-scale algorithmic decision-making. It may also discourage both users and policymakers from taking more robust actions because, after imposing procedural safeguards, they can declare their job done.

Paul Butler made a similar argument about the effect of *Gideon v. Wainwright*¹⁰⁶ on the incarceration of poor persons of color.¹⁰⁷ By focusing on a procedural right to counsel, *Gideon*, Butler argues, obscured the “real crisis of indigent defense” that prison is designed for poor people and not rich ones.¹⁰⁸ Ensuring some adequate representation may not be a bad idea in a vacuum, but it “invests the criminal justice system with a veneer” of legitimacy, impartiality, and protection for ordinary persons, discouraging anyone from digging any deeper into the systematic ways in which the system is stacked against the poor.¹⁰⁹ Butler concluded that, “[o]n its face, the grant that *Gideon* provides poor people seems more than symbolic: it requires states to *pay* for poor people to have lawyers. But the implementation of *Gideon* suggests that the difference between symbolic and material rights might be more apparent than real.”¹¹⁰ Similarly, process-oriented rules for reining in discriminatory algorithms could obscure underlying injustices and stand in the way of substantive reform.

103. See Kenneth A. Bamberger & Deirdre K. Mulligan, *Privacy Decisionmaking in Administrative Agencies*, 75 U. CHI. L. REV. 75, 76 (2008).

104. For Edelman, this is an example of the “mobilization of symbolic structures” to protect the corporation and deny individuals their rights under the law. See EDELMAN, *supra* note 16, at 153–67.

105. Edelman calls this “deference.” See *id.* at 168–71.

106. 372 U.S. 335 (1963).

107. See Paul Butler, *Poor People Lose: Gideon and the Critique of Rights*, 122 YALE L.J. 2176, 2197–98 (2013).

108. *Id.* at 2178.

109. *Id.* at 2178–79.

110. *Id.* at 2191.

C. A Substantive Alternative

Under neoliberal managerialism, which envisions a deregulated market with corporate actors using compliance metrics to largely police themselves,¹¹¹ automated decision-making systems lack democratic legitimacy. They are biased and discriminatory, and legal process is unlikely to address the underlying inequalities and power structures that threaten to make algorithmic systems tools of corporate and majoritarian power.¹¹²

We need a robust, substantive approach to ensure that algorithmic systems meet fundamental social values other than efficiency. To do that, we need to audit the code of automated systems for noncompliance with values like equality, nondiscrimination, dignity, privacy, and human rights. Academic researchers have been doing this for some time and can be deputized to conduct independent sociotechnical analyses of algorithmic systems before and after they are used in commerce or by a government entity.

Consider, for example, the work of Joy Buolamwini and Timnit Gebru and their Gender Shades project.¹¹³ Recognizing the biases in facial recognition technology,¹¹⁴ Buolamwini and Gebru developed an approach to detect, classify, and evaluate the extent of bias in any automated facial analysis program. They deployed their test on three commercially available gender facial recognition systems and found substantial disparities in misclassification among darker females (the most miscategorized), lighter females, darker males, and lighter males (the least miscategorized).¹¹⁵ These systems failed a substantive test for adherence to society's normative values of nondiscrimination and equality and could, therefore, be subject to regulation. Teams at Carnegie Mellon University have determined the extent to which privacy policies are inscrutable and incomprehensible, giving willing regulators ammunition to make notice-and-consent more effective.¹¹⁶

111. See COHEN, *supra* note 16, at 143–59.

112. *Id.* at 4–5 (noting that law and technology are not neutral tools and that both can either be “means of resisting domination or vehicles for embedding it”).

113. See GENDER SHADES, <http://gendershades.org/> [<https://perma.cc/B4LQ-99FC>] (last visited Oct. 6, 2019).

114. See, e.g., Adam Frucci, *HP Face-Tracking Webcams Don't Recognize Black People*, GIZMODO (Dec. 21, 2009, 10:00 AM), <https://gizmodo.com/hp-face-tracking-webcams-dont-recognize-black-people-5431190> [<https://perma.cc/VLQ4-QGTF>]; Loren Grush, *Google Engineer Apologizes After Photos App Tags Two Black People as Gorillas*, VERGE (July 1, 2015, 6:03 PM), <https://www.theverge.com/2015/7/1/8880363/google-apologizes-photos-app-tags-two-black-people-gorillas> [<https://perma.cc/27ZH-XBSF>]; Adam Rose, *Are Face-Detection Cameras Racist?*, TIME (Jan. 22, 2010), <http://content.time.com/time/business/article/0,8599,1954643,00.html> [<https://perma.cc/83XA-5MFW>].

115. See Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, 81 PROC. MACHINE LEARNING 77, 84–87 (2018).

116. See, e.g., Lorrie Faith Cranor, *Necessary but Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice*, 10 J. ON TELECOM. & HIGH TECH. L. 273, 274–75 (2012); Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 I/S: J.L. & POL'Y FOR INFO. SOC'Y. 543, 566–67 (2008). Despite this data, regulators in the United States have generally failed to take action on inconspicuous policies or inscrutable language without any additional corporate mischief.

Princeton University researchers have identified how websites use “dark patterns” to manipulate our behavior,¹¹⁷ offering policymakers the information necessary to trigger regulatory investigations into “unfair and deceptive” business practices.¹¹⁸ Social scientists and researchers at civil society organizations like the AI Now Institute and Data & Society use a variety of tools to analyze the social implications of algorithmic systems.¹¹⁹ Technology companies should be required by law to hire independent technologists, social scientists, ethicists, and other experts as independent ex ante algorithmic auditors, much like the GDPR requires data controllers to hire data protection officers¹²⁰ or how the Sarbanes-Oxley Act requires independent audit committees for public corporations.¹²¹ For ex post analyses, regulators could recruit—and hire to their staffs—similarly diverse experts from computer science, sociology, anthropology, and data science to develop their own substantive tests for algorithmic legitimacy, thereby raising awareness, contributing to necessary changes in design, and creating a robust regulatory agenda.

The tools for this approach to legitimate algorithmic decision-making systems are in place now. Independent academic experts have the expertise and the tools to evaluate automated systems. Sarbanes-Oxley provides a model for independent expertise inside a company. All levels of government, from cities to the federal government, could enact legislation that expresses the values society wants algorithmic decisions to reflect. And regulators can recruit experts to test technological tools in contexts that both protect proprietary information and give outside auditors the data they need. Then we can develop socially conscious algorithmic decision-making systems.

CONCLUSION

The central argument of this Essay is that algorithmic decision-making is a product of the neoliberal managerial project. Neoliberalism’s resistance to social justice and its emphasis on deregulated markets, economic opportunity, and efficiency, coupled with managerialism’s tendency to create

117. See generally Mathur et al., *supra* note 45.

118. See Daniel J. Solove & Woodrow Hartzog, *The FTC and the New Common Law of Privacy*, 114 COLUM. L. REV. 583, 585 (2014) (arguing generally that the FTC’s privacy jurisprudence should be understood as an emerging common law that grows and adapts with new technologies and challenges); see also CHRIS JAY HOOFNAGLE, FEDERAL TRADE COMMISSION: PRIVACY LAW AND POLICY (2016) (describing the origins and multiple ways the FTC protects the privacy of U.S. consumers).

119. See *Initiative: AI on the Ground*, DATA & SOC’Y, <https://datasociety.net/research/ai-on-the-ground/> [<https://perma.cc/B9TR-H6H2>] (last visited Oct. 6, 2019); *Research*, AI NOW INST., <https://ainowinstitute.org/research.html> [<https://perma.cc/3HQS-J5E4>] (last visited Oct. 6, 2019).

120. See GDPR, *supra* note 71, at 7.

121. See Sarbanes-Oxley Act of 2002 § 301, 15 U.S.C. § 78j-1 (2012); Listing Standards Relating to Audit Committees, 17 C.F.R. § 240.10A-3 (2019); Standards Relating to Listed Company Audit Committees, 68 Fed. Reg. 18,788 (Apr. 16, 2003) (to be codified at 17 C.F.R. pts. 228–29, 240, 249, and 274). “Independent” means not being affiliated with the company other than as a director or receiving any compensation other than for serving as a director. See Sarbanes-Oxley Act § 301.

systems that emphasize efficiency and innovation over other values, contribute to a system that values fast, data-driven decision-making by machines. Seen in this way, saying that automated decision-making disempowers humans would be incomplete. It disempowers *some* humans by transforming us into statistics and bits of data. However, algorithmic decision-making empowers engineers and the technology companies they work for because, collectively, they control the social process by which decision-making policy is reformulated into the code that affects whether ordinary individuals receive loans, get put in prison, or gain access to health care. This is the apotheosis of the neoliberal vision. Absent are normative social values like equality, nondiscrimination, and human rights.

This Essay also challenges the conventional wisdom that process and procedure can address the gaps left by automated decision-making. It proposes that regulators, assisted by independent academic experts, audit algorithmic decision-making code for its adherence to social values. And it recommends that tools that fail independent tests should not be deployed.

But in adding to a growing research agenda on algorithms in society, more work needs to be done. Empirical work can determine the bases, if any, on which individuals are willing to accept the decisions of algorithms. Technical research is necessary to develop methodologies for interrogating decision-making code. Interdisciplinary legal research could develop regulatory standards for determining when dark patterns are impermissible.¹²² And legal policy research should determine how best to structure and fund a regulatory body that can do the difficult work of ensuring that algorithmic decision-making systems conform to our social values.

122. Several of these are part of the author's ongoing research, particularly the development of legal standards for regulating dark patterns and empirical work on popular perceptions of the legitimacy of automated decision-making.