

October 2024

## Algorithmic Reason-Giving, Arbitrary and Capricious Review, and the Need for a Clear Normative Baseline

Cameron Averill  
*Yale Law School*

Follow this and additional works at: <https://scholarship.law.uc.edu/uclr>



Part of the [Administrative Law Commons](#), [Artificial Intelligence and Robotics Commons](#), [Science and Technology Law Commons](#), [State and Local Government Law Commons](#), and the [Theory and Algorithms Commons](#)

---

### Recommended Citation

Cameron Averill, *Algorithmic Reason-Giving, Arbitrary and Capricious Review, and the Need for a Clear Normative Baseline*, 93 U. Cin. L. Rev. 40 (2024)

Available at: <https://scholarship.law.uc.edu/uclr/vol93/iss1/2>

This Lead Article is brought to you for free and open access by University of Cincinnati College of Law Scholarship and Publications. It has been accepted for inclusion in University of Cincinnati Law Review by an authorized editor of University of Cincinnati College of Law Scholarship and Publications. For more information, please contact [ronald.jones@uc.edu](mailto:ronald.jones@uc.edu).

ALGORITHMIC REASON-GIVING, ARBITRARY AND  
CAPRICIOUS REVIEW, AND THE NEED FOR A  
CLEAR NORMATIVE BASELINE

*Cameron Averill\**

ABSTRACT

*Federal agencies have caught the artificial intelligence (AI) bug. A December 2023 report by the Government Accountability Office found that twenty of twenty-three federal agencies surveyed reported using some form of AI, with about two hundred current use cases for algorithms and about one thousand more in the planning phase. These agencies are using algorithms in all aspects of administration, including rulemaking, adjudication, and enforcement. The risks of AI are well-documented. Previous work has shown that algorithms can be, among other things, biased and prone to error. However, perhaps no problem poses a more serious threat to the use of algorithms by agencies than the fact that algorithms can be opaque, meaning it can be difficult to understand how an algorithm works and why it reaches certain results. Opacity compromises reason-giving, a basic pillar of administrative governance. Inadequate reason-giving poses legal problems for agencies because the reasons agencies provide for their decisions form the basis of judicial review. Without adequate reason-giving, agency action will fail arbitrary and capricious review under the Administrative Procedure Act. Inadequate reason-giving poses normative problems, too, since reason-giving promotes quality decision making, fosters accountability, and helps agencies respect parties' dignitary interests.*

*This Article considers whether agencies can use algorithms without running afoul of standards, both legal and normative, for reason-giving. It begins by disaggregating algorithmic reason-giving, explaining that algorithmic reason-giving includes both the reasons an agency gives for an algorithm's design (systemic reason-giving) and the reasons an agency gives for an individual decision when the decision making process involves an algorithm (case-specific reason-giving). This Article then evaluates systemic reason-giving and case-specific reason-giving in turn. Once the normative assessment is complete, this Article considers its*

---

\* Yale Law School 2024. I am grateful first and foremost to Abbe R. Gluck and Nicholas R. Parrillo for their involvement in this project since its beginning. For extensive written comments, I thank Alicia Lai and Jerry L. Mashaw. Thank you also to David Freeman Engstrom. Finally, for helpful editing, I thank Madeline Brown as well as Cecilia Giles, Ainsley Ayers, Delaney Cook, Brian Patrick, Audrey Woodward, Christian Thompson, Sebastian Johnson, and the staff of the *University of Cincinnati Law Review*. All errors are my own.

*implications for arbitrary and capricious review, concluding that at least some algorithms should pass judicial muster. The Article finishes by offering a framework that courts can use when evaluating whether the use of an algorithm is arbitrary and capricious, and that agencies can use to decide whether to create an algorithm in the first place.*

*Although understanding the relationship between algorithms and reason-giving is important, this Article's true aim is broader. It seeks to reframe debates over agencies' use of AI by emphasizing that the baseline against which these algorithms should be compared is not some idealized human decision maker, but rather the various kinds of policies—rules, internal procedures, and guidance—that agencies have used since their inception to promote core administrative values like consistency, accuracy, and efficiency. The comparison between algorithms and policies better captures the role algorithms currently play in administrative governance, gives proper weight to the reasons agencies have for turning to algorithms in the first place, and helps us see how algorithms do and do not fit within the existing structures of administrative law. At bottom, comparing algorithms to policies reminds us that the tension between individualized consideration and centralized bureaucratic management is endemic to agency administration. At most, algorithms have given this tension a new flavor. Make no mistake: this tension cannot be eliminated, only managed. Algorithmic reason-giving is a case in point.*

## CONTENTS

|  |    |
|--|----|
| INTRODUCTION .....   | 43 |
| I. ALGORITHMS AS POLICIES .....  | 53 |
| A. <i>Algorithms: A Brief Overview</i> .....                             | 53 |
| B. <i>Agencies' Use of Algorithms</i> .....                              | 54 |
| C. <i>Agencies' Use of Policies</i> .....                                | 59 |
| D. <i>Algorithms as Policies</i> .....                                   | 60 |
| 1. Familiarizing Algorithms: The Medical-Vocational Grid.....            | 61 |
| 2. Breaking it Down: What Algorithms and Policies<br>Have in Common..... | 64 |
| 3. Two Implications for Reason-Giving.....                               | 67 |
| II. SYSTEMIC REASONS FOR ALGORITHMS/POLICIES .....                       | 68 |
| A. <i>What Do Systemic Reasons Look Like?</i> .....                      | 68 |
| B. <i>Virtues of Systemic Reasons for Policies</i> .....                 | 73 |
| 1. Promoting Quality Decision Making .....                               | 73 |
| 2. Fostering Accountability .....  | 74 |
| 3. Preserving Dignitary Interests.....                                   | 74 |
| C. <i>How Systemic Reasons for Algorithms Measure Up</i> .....           | 75 |
| 1. Promoting Quality Decisions.....                                      | 76 |
| 2. Fostering Accountability .....  | 81 |
| 3. Preserving Dignitary Interests.....                                   | 83 |
| III. CASE-SPECIFIC REASONS (OR ALGORITHMS/POLICIES AS REASONS).....      | 84 |
| A. <i>An Old Problem Revisited: Case-Specific Reason-Giving</i> .....    | 85 |
| B. <i>A Way Out: Policies/Algorithms as Reasons</i> .....                | 87 |
| IV. IMPLICATIONS .....   | 92 |
| A. <i>Normative Framework</i> .....                                      | 93 |
| B. <i>Arbitrary and Capricious Review</i> .....                          | 96 |
| CONCLUSION .....   | 99 |

## INTRODUCTION

The Biden Administration has its eyes on artificial intelligence (AI). On October 30, 2023, President Joe Biden issued an Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence.<sup>1</sup> The sprawling order addressed AI's potential effects on, among other things, national security,<sup>2</sup> data privacy,<sup>3</sup> civil rights,<sup>4</sup> consumer protection,<sup>5</sup> labor,<sup>6</sup> and competition.<sup>7</sup> While much of the order concerned private-sector uses of AI, it also instructed the Director of the Office of Management and Budget (OMB) to convene and chair an interagency council to coordinate the development and use of AI by federal agencies.<sup>8</sup> Along those lines, the order outlined the contours of a new bureaucracy consisting of Artificial Intelligence Governance Boards, and Chief Artificial Intelligence Officers within each agency that will be tasked with directing the agency's use of AI.<sup>9</sup> These regulations reflect the Biden Administration's belief that algorithms will be an increasingly important part of administrative governance moving forward.

If current trends are any indication, that belief is well-founded. Federal agencies are already using algorithms in a variety of ways. A December 2023 report by the Government Accountability Office (GAO) found that twenty of twenty-three agencies surveyed reported using some form of AI, with about two hundred current use cases for algorithms and about one thousand more use cases in the planning phase.<sup>10</sup> Similarly, a 2020 report commissioned by the Administrative Conference of the United States (ACUS) found that agencies were using algorithms to support enforcement, adjudication, and rulemaking.<sup>11</sup> These developments are not unique to the federal government, nor to the United States. States and governments around the world are also using automated decision making

---

1. Exec. Order No. 14,110, 88 Fed. Reg. 75,191 (Nov. 1, 2023).

2. *Id.* at 75,196.

3. *Id.* at 75,217.

4. *Id.* at 75,211.

5. *Id.* at 75,214.

6. *Id.* at 75,210.

7. *Id.* at 75,204.

8. *Id.* at 75,218.

9. *Id.*

10. U.S. GOV'T ACCOUNTABILITY OFF., GAO-24-105980, ARTIFICIAL INTELLIGENCE: AGENCIES HAVE BEGUN IMPLEMENTATION BUT NEED TO COMPLETE KEY REQUIREMENTS 17 (2023).

11. See DAVID FREEMAN ENGSTROM, DANIEL E. HO, CATHERINE M. SHARKEY & MARIANO-FLORENTINO CUÉLLAR, GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES (2020) (report to the Admin. Conf. of the U.S.), <https://www.acus.gov/sites/default/files/documents/Government%20by%20Algorithm.pdf> (hereinafter GOVERNMENT BY ALGORITHM).

as part of administration.<sup>12</sup>

This development has been controversial. In the United Kingdom, a grading algorithm led to public outcry.<sup>13</sup> In Arkansas, physically disabled individuals sued the state's Department of Human Services after an algorithm reduced the number of at-home care hours the plaintiffs received by an average of 43%.<sup>14</sup> These episodes remind us that algorithms can create as many problems as they solve. Many of the potential harms of algorithms have been extensively documented. Algorithms can be biased,<sup>15</sup> unpredictable,<sup>16</sup> or riddled with error.<sup>17</sup> Algorithms introduce new security challenges,<sup>18</sup> and privacy risks, too.<sup>19</sup> Even if algorithms are tested, they might struggle to perform in the face of unexpected inputs.<sup>20</sup> Relying on algorithms created by government contractors presents its own set of issues, such as when contractors assert trade secret protection over their algorithms to prevent Due Process challenges.<sup>21</sup>

When it comes to the use of AI by federal agencies, however, one problem predominates: algorithmic opacity. When one says algorithms are opaque, it means that it can be difficult to understand how an

12. See Amos Toh, *The Algorithms Too Few People Are Talking About*, HUMAN RTS. WATCH (Jan. 5, 2024, 11:40 AM), <https://www.hrw.org/news/2024/01/05/algorithms-too-few-people-are-talking-about> (describing the harmful use of algorithms by governments in the United Kingdom, Denmark, Spain, France, and the United States); Eshe Nelson, *European Central Bank Is Experimenting With a New Tool: A.I.*, N.Y. TIMES (Sep. 28, 2023), <https://www.nytimes.com/2023/09/28/business/european-central-bank-artificial-intelligence.html>.

13. See Daan Kolkman, "F\*\*k the Algorithm"?: What the World Can Learn from the UK's A-Level Grading Fiasco, LSE: IMPACT OF SOC. SCIS. (Aug. 26, 2020), <https://blogs.lse.ac.uk/impactofsocialsciences/2020/08/26/fk-the-algorithm-what-the-world-can-learn-from-the-uks-a-level-grading-fiasco/>.

14. See *infra* Section I.B.

15. See VIRGINIA EUBANKS, *AUTOMATING INEQUALITY: HOW HIGH-TECH TOOLS PROFILE, POLICE, AND PUNISH THE POOR* (2017); SAFIYA UMOJA NOBLE, *ALGORITHMS OF OPPRESSION: HOW SEARCH ENGINES REINFORCE RACISM* (2018); CATHY O'NEIL, *WEAPONS OF MATH DESTRUCTION: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY* (1st ed. 2016).

16. See Andrew Smith, *Franken-Algorithms: The Deadly Consequences of Unpredictable Code*, THE GUARDIAN (Aug. 30, 2018, 01:00 AM), <https://www.theguardian.com/technology/2018/aug/29/coding-algorithms-frankenalgos-program-danger>.

17. See Mike Ananny, *Seeing Like an Algorithmic Error: What are Algorithmic Mistakes, Why Do They Matter, How Might They Be Public Problems?*, 24 YALE J.L. & TECH. 342, 347 (2022).

18. See, e.g., Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. DAVIS L. REV. 399, 419-20 (2017) (explaining cybersecurity challenges created or exacerbated by AI).

19. See, e.g., Karl Manheim & Lyric Kaplan, *Artificial Intelligence: Risks to Privacy and Democracy*, 21 YALE J.L. & TECH. 106, 116-33 (2019) (detailing privacy risks of AI); Alicia Solow-Niederman, *Information Privacy and the Inference Economy*, 117 NW. U.L. REV. 357, 378 (2022) (explaining new privacy risks created by machine learning algorithms).

20. See *infra* Section II.C.1.

21. See Cary Coglianese, *AI, Due Process, and Trade Secrets*, REG. REV. (Sept. 4, 2023), <https://www.thereview.org/2023/09/04/coglianese-ai-due-process-and-trade-secrets/>.

algorithm reaches a result.<sup>22</sup> “Opacity” here captures two problems.<sup>23</sup> First, algorithms can be opaque because understanding how they work requires technical expertise.<sup>24</sup> For example, the Social Security Administration (SSA)—a leader among federal agencies in creating algorithms—has hired engineers who translate agency priorities into code.<sup>25</sup> Even if these engineers understand how an algorithm reaches a result, the people using the algorithm to make decisions might not. Likewise, the people affected by the algorithm—say, disability claimants—might also lack necessary technical expertise to understand the algorithms.

Second, algorithms can be opaque when it is not possible *for anyone* to explain why the algorithm produced a result in a particular case. This is often called the black-box problem.<sup>26</sup> It is a particular concern for machine learning algorithms, which rely on statistics and large datasets to reach their results.<sup>27</sup> An algorithm might be a black box when it correlates variables in a non-intuitive way,<sup>28</sup> or when it produces outputs based on more factors than a human being considers when making decisions. The black-box problem has already caught the attention of federal agencies worried about algorithms being used in the private sector.<sup>29</sup> This Article considers the significance of algorithmic opacity for the algorithms that agencies are using themselves.

22. See Jenna Burrell, *How the Machine ‘Thinks’: Understanding Opacity in Machine Learning Algorithms*, *BIG DATA & SOC’Y*, June 2016, at 1-2.

23. This Article does not address opacity resulting from deliberate secrecy. For that topic, see, for example, Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 *FORDHAM L. REV.* 678, 1092 (2018), which explains why keeping algorithms secret is problematic. This kind of opacity is especially prevalent when agencies use algorithms built by third-party contractors. This Article generally posits that agencies’ use of algorithms is legitimate to the extent that agencies give reasons for an algorithm’s design. Suffice it to say, then, that under this Article’s framework, agencies’ use of trade-secret algorithms built by contractors is suspect.

24. See W. Nicholson Price II & Arti K. Rai, *Clearing Opacity through Machine Learning*, 106 *IOWA L. REV.* 775, 784 (2021).

25. See *infra* Section II.C.1.

26. See Yavar Bathaee, *The Artificial Intelligence Black Box and the Failure of Intent and Causation*, 31 *HARV. J.L. & TECH.* 889, 905 (2018) (“Generally, the Black Box Problem can be defined as an inability to fully understand an AI’s decision making process and the inability to predict the AI’s decisions or outputs.”).

27. See *infra* Section I.A.

28. See Selbst & Barocas, *supra* note 23, at 1097.

29. See, e.g., Michael Atleson, *Keep Your AI Claims in Check*, FTC (Feb. 27, 2023), <https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check> (“[Y]ou can’t say you’re not responsible because that technology is a ‘black box’ you can’t understand or didn’t know how to test.”); *CFPB Acts to Protect the Public from Black-Box Credit Models Using Complex Algorithms*, CFPB (May 26, 2022), <https://www.consumerfinance.gov/about-us/newsroom/cfpb-acts-to-protect-the-public-from-black-box-credit-models-using-complex-algorithms/> (“Whistleblowers play a central role in uncovering information about companies using technologies, like black-box models, in ways that violate ECOA and other federal consumer financial protection laws.”).

Opacity is one of the most urgent problems relating to federal agencies' use of algorithms. No doubt, it has competition. Bias, error, threats to privacy, and other problems are serious, concerning, and worthy of attention. However, unlike these other issues, opacity affects *all* algorithms to some degree. More importantly, opacity strikes at a core pillar of administration: reason-giving.

Reason-giving is central to administrative law.<sup>30</sup> The Administrative Procedure Act (APA) requires agencies to give reasons for all agency action.<sup>31</sup> Without adequate reason-giving, agency action will fail arbitrary and capricious review, under which courts must “hold unlawful and set aside” agency action that is “arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with law.”<sup>32</sup> But the problem runs deeper than that. The reasons an agency gives form the record for judicial review under *any* standard—Due Process, substantial evidence, and so on. By law, a reviewing court can only uphold agency action based on the original reasons the agency gave for acting—a requirement older than the APA itself.<sup>33</sup> Legality aside, there are good reasons for requiring reasons. Recent experiments by Edward Stiglitz indicate that reason-giving “constrains agency action and encourages other regard and attention to statutory objectives.”<sup>34</sup> Requiring agencies to give reasons for their decisions also can improve decisional accuracy, promote efficiency, strengthen agency legitimacy, and foster accountability.<sup>35</sup>

Algorithmic opacity undermines reason-giving in at least two ways, corresponding to the two kinds of opacity mentioned above. First, agency employees making decisions based on algorithms might lack the expertise needed to understand the technical reasons given for an algorithm's design. The individuals subject to regulation might lack such expertise, too. Second, there is the problem of the algorithmic black box. One usually cannot ask a machine learning algorithm why it produced a result in a particular case. Instead, one is consigned to understanding how the

---

30. This Article owes a large debt, in structure and content, to Ashley S. Deeks, *Secret Reason-Giving*, 129 YALE L.J. 612 (2020), which describes reason-giving the Executive undertakes in secret and argues that such reason-giving should become a regular part of the Executive's decision making process. For more on the importance of reason-giving to administrative governance, see EDWARD H. STIGLITZ, *THE REASONING STATE* 138 (2022); JERRY L. MASHAW, *REASONED ADMINISTRATION AND DEMOCRATIC LEGITIMACY* 1-12 (2018).

31. See 5 U.S.C. § 557(e)(3) (2018) (“[T]he parties are entitled to a reasonable opportunity to submit for the consideration of the employees participating in the decisions . . . supporting reasons for the exceptions or proposed findings or conclusions.”).

32. 5 U.S.C. § 706 (2018).

33. *SEC v. Chenery Corp.*, 318 U.S. 80, 95 (1943).

34. See Stiglitz, *supra* note 30, at 175.

35. See Deeks, *supra* note 30, at 626-34 (articulating virtues of reason-giving).



algorithm works in other ways, such as auditing.<sup>36</sup> That could lead to a situation in which an agency employee—say, a frontline adjudicator making a sensitive benefits decision—does something simply “because the algorithm said so.”

Existing work on algorithmic opacity has followed two general approaches.<sup>37</sup> The first approach emphasizes the ways that different modes of explanation, including audits, data descriptions, and system-level overviews, can make algorithms more comprehensible, while the second calls for simplifying the algorithms themselves to make them more understandable.<sup>38</sup> This Article largely takes the first approach, but it concedes that using algorithms may lead agencies to give worse reasons for their individual decisions than if they relied entirely on human judgment. No matter, this Article contends: agencies should, under certain conditions, be able to use such algorithms regardless.

Ultimately, this Article asks whether “because the algorithm said so” can ever be enough of a reason to justify agency action, from both a legal and a normative perspective.<sup>39</sup> This Article answers yes, but it is important to be precise about why. This Article does not reach that answer through techno-optimism. Instead, it gets there by considering how algorithms compare to the best alternative. What we take that reasonable alternative to be makes all the difference.

Briefly, we must put reason-giving on the back burner and ask what role algorithms play within administrative agencies. As an answer, this Article advances a broad theoretical argument: agency algorithms are like *policies*, not people. Algorithms resemble policies in their goals, in how they are created, and in their effects on human decision making. Both are tools through which agency management can standardize and centralize decision making. Neither removes humans from the decision making process. Instead, algorithms operate in conjunction with humans, shifting some decision making away authority from frontline adjudicators in the process. To create algorithms and policies, agencies invest significant time, energy, and thought on the front end. They do so in the hopes that the algorithm or policy will encode the expertise of many individuals, disseminating it uniformly across the organization.

Consider an example. The Food and Drug Administration (FDA) has

---

36. See Aziz Z. Huq, *Constitutional Rights in the Machine-Learning State*, 105 CORNELL L. REV. 1875, 1944-48 (2020) (discussing methods for making algorithms more scrutable).

37. See GOVERNMENT BY ALGORITHM, *supra* note 11, at 75 (providing this taxonomy).

38. *Id.*; Selbst & Barocas, *supra* note 23, at 1109-17 (detailing techniques for explaining machine learning models); 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, 55-59 (2017) (cataloguing modes of explanation).

39. For an approach that focuses on the legal question, see Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 ADMIN. L. REV. 1, 45-47 (2019).

piloted algorithms to help the agency conduct post-market surveillance of drugs. Under the Food and Drug Administration Amendments Act of 2007, the FDA has the authority to “require a drug sponsor to conduct post-approval studies or new clinical trials at any time after approval of a new drug application if [the] FDA becomes aware of new safety information . . . to require labeling changes to disclose new safety information . . . and to require ‘risk evaluation and management strategies[.]’”<sup>40</sup> To identify problems with drugs on the market, the FDA solicits adverse event reports from patients, caregivers, and especially manufacturers, who are required to submit reports.<sup>41</sup> Since at least 2017, the FDA has piloted machine learning tools that can help sort through these troves of unstructured data.<sup>42</sup> One tool helped the agency sort reports by their likelihood of containing information about severe events.<sup>43</sup> While employees still read all reports, the algorithm helps identify the most pressing ones. The algorithm thus acts like a triaging policy to help the agency make informed judgments with limited resources.<sup>44</sup>

The algorithms-as-policies view stands in contrast to existing work evaluating algorithms against human decision makers. In their impressive defense of algorithms in administrative governance, Cary Coglianese and Alicia Lai argue that algorithms are often more comprehensible than humans.<sup>45</sup> These authors note that “[a]ny meaningful assessment of AI in the public sector must . . . start with an acknowledgment that government as it exists today is already grounded in a set of imperfect algorithms. These existing algorithms are inherent in human decision making.”<sup>46</sup> They conclude that “to the extent that automated systems based on digital algorithms would make improvements over human algorithms for specific tasks, they should be adopted.”<sup>47</sup>

Coglianese and Lai are correct that we should not demand perfect algorithms, since the status quo is far from it. But the important question, is what the baseline for comparison—the assumed status quo—should be. Contrasting algorithms with human judgment can be fruitful. But, this Article argues, it is not the *most* illuminating possible comparison. It is better to view algorithms not as substitute humans but as tools of

---

40. See 21 U.S.C. § 355 (2022).

41. See GOVERNMENT BY ALGORITHM, *supra* note 11, at 55.

42. *Id.* at 53.

43. *Id.* at 55.

44. See *id.* (“Much like the SEC’s enforcement tools or the SSA’s case clustering tool profiled above, the tool can be thought of as performing a kind of triage to better target scarce agency resources rather than displacing human assessments.”).

45. Cary Coglianese & Alicia Lai, *Algorithm vs. Algorithm*, 71 DUKE L.J. 1281, 1287 (2022).

46. *Id.* at 1286.

47. *Id.* at 1287.

bureaucratic administration—as, to use an umbrella term, a kind of policy.<sup>48</sup> For one, that is because algorithms do not usually eliminate human judgment—they work with it.<sup>49</sup> For another, it is because agency policies already constrain human judgment to varying degrees. The human-versus-algorithms comparison can make it hard to see which problems are distinctive to algorithms and which are inherent to administration.

In this Article’s parlance, a “policy” is a rule, formulated in advance of a particular decision by someone other than the decision maker, that directs an outcome for different possible states of the world.<sup>50</sup> By using the term “rule,” this Article does not mean to contrast policies with standards. In this Article’s terminology, “policies” encompasses rules, standards, and guidelines of all kinds.<sup>51</sup> Some policies, called legislative

---

48. This Article is not the first to compare algorithms to policies, even if other authors often leave the comparison implicit. One piece that deserves special note is Danielle Citron’s pathbreaking article, *Technological Due Process*. Citron recognized that algorithms sometimes play a similar role to rules that apply across adjudications. See Danielle Keats Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1278 (2008). The canonical example of such a rule, described in depth below, is the medical-vocational grid used by the Social Security Administration to adjudicate disability claims, which was made famous in *Heckler v. Campbell*. See *infra* Section I.D.1. Citron’s piece shows why comparing algorithms to policies is fruitful. For instance, the comparison leads her to recognize that algorithms could be outcome determinative in adjudications but, in contrast to rules like the medical-vocational grid, not go through notice and comment. But more than past work, this Article makes an explicit argument for why the algorithm-policy comparison should be the starting point when evaluating agencies’ use of algorithms. Additionally, this Article emphasizes that algorithms are not necessarily replacing rulemaking in all contexts. That is why this Article uses the term “policy” rather than rule—to capture the wide range of bureaucratic tools agencies use to standardize decision making, including legislative rules, internal procedures, and guidance documents. This difference in terminology reflects the fact that policies are endemic to agency adjudication. These policies need not take the form of formal rules. Citron is right that algorithms straddle the line between adjudication and rulemaking. This Article just questions how unique such straddling is to algorithms. *Id.* at 1278. Some courts have also considered, albeit obliquely, where algorithms fit within bureaucratic administration. In *NRDC v. EPA*, for example, the court held that a model used to forecast the likely responses of automakers to proposed emissions standards was not “deliberative” and thus did not fall within the scope of the deliberative process privilege. 954 F.3d 150, 159 (2d Cir. 2020). In the court’s view, the model did not reflect policy judgments to such a degree that it fell within the privilege. Instead, the court viewed the model at issue as a kind of calculator—its use so routine that it did not significantly implicate agency discretion. *Id.* at 157. The calculator view of algorithms risks understating the extent to which algorithms shape and reflect policy judgments. The *NRDC* court seemed to recognize as much and clarified that some algorithms could fall within the deliberative process privilege. *Id.* at 158 n.7.

49. Rebecca Crotoof, Margot E. Kaminski & W. Nicholson Price II, *Humans in the Loop*, 76 VAND. L. REV. 429, 443 (2023).

50. This definition is inspired by the one offered in Nabil I. Al-Najjar, *A Bayesian Framework for the Precautionary Principle*, 44 J. OF LEGAL STUD. S337, S341 (2015) (“A policy is, formally, an act that indicates which consequence obtains at each state of the world. The adoption of a medical treatment, in this language, is an act *f* that yields consequence *f*(*s*) when the state happens to be *s*. Approving an alternative treatment (or doing nothing and maintaining the status quo) corresponds to the selection of another act *g*.”).

51. This Article leaves it to future work to determine whether the algorithms used by agencies are most like rules, standards, or something else entirely. An algorithm resembles a rule insofar as it, to borrow Danielle Citron’s language, “prescribes ex ante an outcome for a particular fact scenario.” See Citron,

rules, are vetted thoroughly before adoption through notice and comment rulemaking. For other policies, such as guidance documents, the reason-giving process is less formal, less public, and may involve less involvement from stakeholders outside of the agency. Some policies apply to individuals and entities outside of the agency, while other policies are for internal administration.

The central point of the algorithms-as-policies framework is this: when evaluating agencies' use of algorithms, the proper picture against which to compare algorithms is not a fictionalized environment in which an individual bureaucrat—or, more appropriately, thousands of bureaucrats—makes decisions with unfettered discretion. Rather, the picture of the administrative state we should have in mind is one in which policies constantly mediate, and sometimes even dictate, human decision making. Agencies could not function without policies. That is not to say agencies ought to involve policies—or, for that matter, algorithms—in every administrative decision. Speaking generally, however, agencies need policies to structure administrator decision making. The critical normative question is whether and when they should use algorithms instead. This Article addresses that question by assessing the implications of the algorithms-as-policies framework for reason-giving. It focuses on reason-giving because reason-giving sits at the heart of administrative law. If administrative law tends to channel complex substantive disputes into narrow procedural capillaries, reasons are the blood cells that flow through.<sup>52</sup> If algorithms create normative and legal problems for reason-giving, algorithms may be unsuitable for administrative governance.

To assess the suitability of algorithms for administration, this Article starts by noting that reason-giving encompasses two different processes. First, it includes giving reasons for designing an algorithm or policy in a particular way. This Article calls these reasons “systemic reasons” because they relate to the design of the overall system. But reason-giving also entails giving reasons in a particular case *based on* an algorithm or policy. These are “case-specific reasons.” This Article considers systemic reasons first. Like all reasons, systemic reasons ought to promote quality decision making, foster accountability, and preserve dignitary interests. This Article explains how systemic reasons for policies advance these interests, and then argues that systemic reasons for algorithms can adequately promote them, too. That said, algorithms do pose one

---

*supra* note 48, at 1301. However, machine learning algorithms challenge this picture because they may be trained on past agency decisions for which standards and discretion predominated. Further complicating the question is that algorithms are rarely dispositive in the federal agency context, at least currently. Human discretion is usually still present to at least some degree.

52. JERRY L. MASHAW, BUREAUCRATIC JUSTICE: MANAGING SOCIAL SECURITY DISABILITY CLAIMS 19 (1983).

important, distinctive problem for systemic reason-giving—namely, that some systemic reasons for algorithms are necessarily technically complex. While this deficiency raises quality, accountability, and dignity concerns, none are so great as to justify excluding algorithms wholesale, especially when these algorithms are compared to similarly complex policies.

Next, this Article turns to case-specific reasons. It begins by explaining that both algorithms and policies undermine individualized reason-giving. Using controversy over the medical-vocational guidelines at issue in *Heckler v. Campbell*, this Article explains how policies, like algorithms, detach reason-giving from the facts of a specific case.<sup>53</sup> Thus, the reason for a particular decision often amounts to little more than, “because the algorithm/policy says so.” Both algorithms and policies thereby privilege systemic management over individualized consideration. How much an agency should prioritize systemic decision making versus individualized consideration is a core problem of administration. This Article argues that however one may distinguish between acceptable and unacceptable limits on individualized consideration—and one must draw this line somewhere—the line for policies and the line for algorithms should be drawn relatively close together.

Still, it is worth asking why algorithms and policies are justifiable from a reason-giving perspective—why, in other words, it can ever be enough to say, “because the algorithm/policy said so.” The key is to understand that algorithms and policies can be reasons for agency action in and of themselves. Drawing on the work of Joseph Raz, and more recent work by Blake Emerson, this Article explains that policies are reasons in two ways.<sup>54</sup> First, they are reasons to take some certain action. For instance, when a police officer pulls over a driver for speeding, a speed limit can, in and of itself, be a reason for writing the driver a ticket. Second, a policy is also a reason *not* to act for competing reasons. In other words, policies narrow the field of what must be considered. In this sense, policies are exclusionary. The black-box problem is less vexing if an algorithm is a valid reason for agency action, since it means “because the algorithm said so” is an adequate reason for an agency to reach a particular decision.

Like policies, algorithms can be valid reasons for an agency to act. Primarily, policies can be reasons for agency action for reasons of necessity. Although reason-giving is more individualized in their absence, policies are what allow agencies to achieve consistency, accuracy, and efficiency goals. If algorithms further these same goals, they should be able to function as reasons, too. This Article argues as much, but with the

---

53. See *Heckler v. Campbell*, 461 U.S. 458, 459-60 (1983).

54. See Joseph Raz, *Reasoning with Rules*, in *BETWEEN AUTHORITY AND INTERPRETATION: ON THE THEORY OF LAW AND PRACTICAL REASON* 216 (Joseph Raz ed., 2009).

caveat that agencies should try to make public—at least to the relevant stakeholders—the systemic reasons for the algorithm, as these are the evidence that administrative goals are being met. Without systemic reasons, an algorithm is probably not, in and of itself, a reason for agency action.<sup>55</sup>

Algorithms are not perfect as reasons for agency action. Compared to many policies, they are less intuitive on their face, which raises dignitary, accountability, and accuracy concerns. That intuition gap, along with the inherent technical complexity of some systemic reasons for algorithms, means that algorithms may be worse for reason-giving than policies, but still good enough to play a role in the administrative state. Algorithmic reason-giving is not so inferior that reason-giving alone should dictate whether an agency uses an algorithm or policy to address a particular need. Nor is algorithmic reason-giving so flawed as to make using algorithms per se arbitrary and capricious. Both agencies and courts should consider whether potential costs to reason-giving are outweighed by other benefits such as improvements in accuracy, efficiency, and consistency. Accordingly, this Article offers a framework that can help agencies and courts decide, based mainly on the systemic reasons given for an algorithm, whether the agency’s use of the algorithm is defensible.

This Article proceeds as follows: Section I lays out the algorithms-as-policies framework. Section II applies this framework to systemic reason-giving. Section III considers algorithms and policies as reasons for agency action. Section IV finishes by offering the normative framework and asking what this entire discussion means for arbitrary and capricious review.

Discussions tying algorithms to administrative law are rare, as is work that asks whether agencies’ use of algorithms is normatively defensible.<sup>56</sup>

---

55. While Section IV.A of this Article introduces a normative framework for determining whether the reasons given for an algorithm are adequate, the framework assumes that at least some reason-giving process is taking place. It is outside the scope of this Article to explain the ideal process for developing an agency algorithm. A process like notice and comment rulemaking may foster the greatest volume of systemic reason-giving and may invite the most public involvement, while an entirely internal development process might lead an agency to offer fewer, and potentially lower quality, systemic reasons for an algorithm. At this stage, it at least seems correct to say that for algorithms to be equally legitimate to policies as reasons for agency action, algorithms and policies serving similar functions within an agency should undergo similarly rigorous and predictable processes before adoption. For an effort to adapt the APA’s doctrinal categories to agency algorithms, see Peter Henderson & Mark Krass, *Algorithmic Rulemaking vs. Algorithmic Guidance*, 37 HARV. J.L. & TECH. 105 (2013). Thank you to Alicia Lai for raising this point about process. For a further discussion of process, including a discussion of what Coglianese and Lai identify as the “meta-process,” see *infra* Section IV.A.

56. See David Freeman Engstrom & Daniel E Ho, *Algorithmic Accountability in the Administrative State*, 37 YALE J. REG. 800, 805 (2020) (“Only a trickle of research treats the more fine-grained statutory requirements of administrative law and, even then, offers mostly a surface-level tour of potentially applicable doctrines.”). Notable exceptions include Danielle Citron & Ryan Calo, *The Automated Administrative State: A Crisis of Legitimacy*, 70 EMORY L.J. 797 (2021); Mariano Florentino

This Article tries to start filling the gap with a consideration of algorithmic reason-giving. The broader hope is that the algorithms-as-policies framework will prove helpful to scholars addressing other problems related to agencies' use of AI.

### I. ALGORITHMS AS POLICIES

What role do algorithms play in administration? This Section addresses that question by considering the algorithms that federal agencies are building in-house for use in adjudication, enforcement, and rulemaking. It begins by explaining how machine learning algorithms are created. It then describes how agencies are currently using algorithms. After a brief discussion of how agencies use policies to structure decision making, this Article explains why comparing algorithms and policies is worthwhile.

#### A. Algorithms: A Brief Overview

Broadly speaking, computerized algorithms fall into two categories: rules-based and machine learning. Rules-based algorithms, as the name implies, rely on rules explicitly hard-coded by humans. The simplest version of a rules-based algorithm would be something like, "if X, do Y." Machine learning algorithms use rules, too, but humans do not hard-code them. Instead, the algorithm "learns" rules by performing statistical calculations of various kinds on large datasets. Humans play a large role in creating machine learning algorithms; they just do not specify explicit rules for weighing inputs. This Article mostly focuses on machine learning algorithms because, as explained below, rules-based algorithms are essentially indistinguishable from more traditional policies and therefore do not raise similar concerns about reason-giving.<sup>57</sup>

To create a machine learning algorithm, members of the design team take the following steps.<sup>58</sup> First, they decide what function they want the algorithm to perform and choose an outcome variable to be predicted by the algorithm.<sup>59</sup> Second, they collect the data.<sup>60</sup> This step is consequential. The accuracy of a machine learning algorithm hinges in part on the data

---

Cuéllar, *Cyberdelegation and the Administrative State*, in ADMINISTRATIVE LAW FROM THE INSIDE OUT: ESSAYS ON THEMES IN THE WORK OF JERRY L. MASHAW 134 (Nicholas R. Parrillo ed., 2017); Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147 (2017); Deirdre K. Mulligan & Kenneth A. Bamberger, *Procurement as Policy: Administrative Process for Machine Learning*, 34 BERKELEY TECH. L.J. 773 (2019).

57. See *infra* Section I.D.1.

58. These steps are outlined in fuller detail in David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653 (2017).

59. *Id.* at 673-74.

60. *Id.* at 677.

used to train it. The volume of data is important; in some domains, the best-performing algorithms rely on datasets of hundreds of millions of entries.<sup>61</sup> To create datasets, many teams use low-wage workers to tag, or label, data points with different features.<sup>62</sup> These workers also help with the third step, which is to clean the data.<sup>63</sup> Cleaning the dataset means removing errors that may have come up during its compilation. Fourth, the team may run summary statistics to get a sense of overarching patterns in the data.<sup>64</sup> Fifth, the team partitions the dataset. In this step, the team must sequester some of the data into a “test” dataset—in contrast to the “training” dataset—which the team uses to measure the algorithm’s accuracy and performance.<sup>65</sup> Sixth, the team chooses what kind of machine learning model to use.<sup>66</sup> Machine learning algorithms come in a variety of forms and can use different mechanisms for predictions.<sup>67</sup> Seventh, the team runs the chosen method on the training dataset.<sup>68</sup> This process is iterative and involves tuning, assessment, and feature selection (choosing or excluding inputs to go into the algorithm).<sup>69</sup> A great degree of trial-and-error is often necessary to achieve desired levels of accuracy.<sup>70</sup> Finally, the team can deploy the algorithm.<sup>71</sup> Once this is done, the team ideally will engage in ongoing auditing to see how the algorithm performs “in the wild.”

### *B. Agencies’ Use of Algorithms*

State and federal agencies are using algorithms to perform a variety of tasks. Thanks to recent work by scholars working for ACUS, it is not necessary to rely entirely on hypothetical examples to examine the problems posed by the government’s use of algorithms.<sup>72</sup> Although general language models like ChatGPT have recently garnered headlines,

---

61. See, e.g., *Datasets, Generalization, and Overfitting*, GOOGLE DEVELOPERS, <https://developers.google.com/machine-learning/data-prep/construct/collect/data-size-quality> (giving the size of datasets used to train popular Google products).

62. See KATE CRAWFORD, *THE ATLAS OF AI: POWER, POLITICS, AND THE PLANETARY COSTS OF ARTIFICIAL INTELLIGENCE* 63-64 (2021).

63. See Lehr & Ohm, *supra* note 58, at 681.

64. *Id.* at 683.

65. *Id.* at 685.

66. *Id.* at 688.

67. *Id.* at 689.

68. *Id.* at 695.

69. *Id.* at 696.

70. *Id.* at 698.

71. *Id.* at 701.

72. *Artificial Intelligence*, ADMIN. CONF. OF THE U.S., <https://www.acus.gov/ai> (2023) (“This page provides access to many of the publications, projects, and other materials related to AI that have been prepared by ACUS or for its consideration. It will be updated as new materials become available.”).



the algorithms agencies have built thus far have performed specific tasks in only specific domains. Public-facing adjudicatory tools at the state level have sparked the most controversy, but at least at the federal level, a large portion of the algorithms that agencies are using perform internal, bureaucratic, and relatively mundane tasks. Algorithms binding final agency action are, at this point, rare, and possibly nonexistent. That said, algorithms are not isolated to any single use. At the federal level, agencies are using algorithms in all realms of agency action, including adjudication, rulemaking, and enforcement.<sup>73</sup>

In adjudications, algorithms perform important functions. At the federal level, however, they are not yet outcome determinative. One example of an algorithm used in adjudication is the SSA's Quick Disability Determination (QDD) algorithm, which is used to expedite certain applications for benefits.<sup>74</sup> QDD uses a predictive machine learning model to identify claimants who are likely to prevail on their disability claims. As a claims representative evaluates an application, the QDD model scores the case and alerts the claims representative if the case qualifies as a QDD case. Adjudicators making the final decision do not have access to the predicted probability of success, likely out of concern that the score will influence their likelihood of granting benefits. Still, QDD has important effects on regulated parties. For one, it determines who gets benefits the fastest. For another, QDD cases are assigned to designated examiners specifically trained to conduct QDD adjudications.

Other examples abound. For instance, the United States Patent and Trademark Office uses an algorithm to predict design codes to assign to new trademarks.<sup>75</sup> It uses another to search through existing trademarks.<sup>76</sup> The SSA has also piloted algorithms to identify anomalies in hearing decisions, with mixed results.<sup>77</sup>

State agencies have used adjudicatory algorithms more aggressively, sometimes with devastating results.<sup>78</sup> In one especially concerning case, Arkansas's Department of Human Services built algorithms to determine

---

73. For a catalogue of algorithms used by federal agencies, see GOVERNMENT BY ALGORITHM, *supra* note 11, appendix.

74. See *Quick Disability Determinations*, SOC. SEC. ADMIN., <https://www.ssa.gov/disabilityresearch/qdd.htm>.

75. See *Emerging Technologies in USPTO Business Solutions*, WORLD INTELL. PROP. ORG. (2018), [https://www.wipo.int/edocs/mdocs/globalinfra/en/wipo\\_ip\\_itai\\_ge\\_18/wipo\\_ip\\_itai\\_ge\\_18\\_p5.pdf](https://www.wipo.int/edocs/mdocs/globalinfra/en/wipo_ip_itai_ge_18/wipo_ip_itai_ge_18_p5.pdf) (slide 14).

76. *Id.* (slide 19).

77. *The Social Security Administration's Use of Insight Software to Identify Potential Anomalies in Hearing Decisions*, OFF. OF THE INSPECTOR GEN. (2019).

78. See, e.g., Citron & Calo, *supra* note 56, at 818-32 (compiling examples from litigation); TRANSFORMING DELIVERY OF HEALTH & HUMAN SERVICES THROUGH ROBOTICS PROCESS AUTOMATION (2019), [https://www.nascio.org/wp-content/uploads/2020/09/NASCIO-Awards-2019\\_State-of-OH-Bots.pdf](https://www.nascio.org/wp-content/uploads/2020/09/NASCIO-Awards-2019_State-of-OH-Bots.pdf) (describing automated tools for administering public assistance programs).

how many at-home care hours low-income individuals with physical disabilities would receive to complete daily tasks such as bathing, eating, and going to the bathroom.<sup>79</sup> Previously, individual nurses had consulted with patients to make these decisions. The algorithms, however, removed any discretion from the individual nurses, and nurses could not override the algorithms' determinations as to at-home care hours.<sup>80</sup> Plaintiffs filed suit to challenge the validity of the algorithms, alleging that the switch to an algorithm reduced their care hours by an average of 43%.<sup>81</sup> Notably, the algorithm was a rules-based—not a machine learning—algorithm, meaning humans explicitly signed off on the rules governing decisions.<sup>82</sup> Also critically, the algorithm was built by a third party and never exposed to public scrutiny, making it harder to uncover errors.<sup>83</sup>

Agencies are also using algorithms in the rulemaking process. For example, agencies currently use machine learning algorithms to sort through and categorize submissions during notice and comment.<sup>84</sup> While most rulemaking involves only a limited set of stakeholders,<sup>85</sup> agencies can receive hundreds of thousands of comments for a single proposed rule. In the past, this has led agencies to outsource some of the work of reading and grouping comments to third-party contractors.<sup>86</sup> But now, algorithms do some of the up-front categorization work.<sup>87</sup> Agencies have also piloted algorithms that can assist agency employees in developing more informed rules, as the FDA example discussed in Section I shows.

Although such an algorithm is not yet in place, one could imagine an algorithm playing an even greater role than the QDD in determining what procedures are mandatory during rulemaking. Although the APA requires

79. Ark. Dep't of Hum. Servs. v. Ledgerwood, 530 S.W.3d 336, 339 (Ark. 2017).

80. *Id.*

81. *Id.*

82. Colin Lecher, *A Healthcare Algorithm Started Cutting Care, and No One Knew Why*, THE VERGE (Mar. 21, 2018), <https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy> (“[T]he algorithm computes about [sixty] descriptions, symptoms, and ailments — fever, weight loss, ventilator use — into categories, each one corresponding to a number of hours of home care.”).

83. Ark. Dep't of Hum. Servs. v. Ledgerwood, 530 S.W.3d at 345.

84. Bridget C.E. Dooling & Rachel Augustine Potter, *Contractors in Rulemaking* 28 (2022), <https://www.acus.gov/sites/default/files/documents/Contractors%20in%20Rulemaking%20Draft%20Report%202022.03.17.pdf> (“This could also include the use of natural language processing tools to help identify mass comment campaigns, for example, or subject matter themes.”).

85. Brian D. Libgober, *Strategic Proposals, Endogenous Comments, and Bias in Rulemaking*, 82 J. OF POL. 642, 642 (2020) (“[R]ulemaking usually involves conflict between a more limited set of stakeholders.”).

86. See Dooling & Potter, *supra* note 84, at 28.

87. It is not clear how many agencies use such tools in notice and comment. While at least some do, scholars have encouraged more agencies to use algorithms in the notice and comment stage to promote transparency and facilitate public participation. See GOVERNMENT BY ALGORITHM, *supra* note 11, at 64.

rulemaking to follow several steps,<sup>88</sup> some agencies go beyond these requirements. For instance, the Department of Transportation (DOT) has policies that require the agency to implement enhanced procedures during rulemaking for rules deemed likely to be “economically significant.”<sup>89</sup> For such rules, the notice of proposed rulemaking must include a discussion explaining an achievable objective, the notice and comment period must be at least sixty days (or ninety days for high impact rules), and any interested party can petition for a formal hearing of the proposed rule according to procedures set forth in the APA.<sup>90</sup> While employees have discretion to deny formal hearings, the DOT has an internal policy setting forth criteria under which an agency employee can deny a petition for a formal hearing.<sup>91</sup>

The DOT is not alone in imposing heightened procedural requirements on significant rulemakings. Under Executive Order No. 12866, issued during the Clinton Administration, all executive agencies must conduct cost-benefit analyses for major rulemakings.<sup>92</sup> However, Executive Order No. 12866 only applies to executive agencies.<sup>93</sup> Executive Order No. 13579, by contrast, encourages independent agencies, to the extent permitted by law, to consider the costs and benefits of rules.<sup>94</sup> The Federal Deposit Insurance Corporation (FDIC), an independent agency, conducted an investigation in 2020 to determine whether the agency was complying with best practices for cost-benefit analysis, as described by the OMB, the GAO, the ACUS, other agencies, and academics.<sup>95</sup> It found that it lacked a documented process for determining when and how to conduct cost-benefit analyses.<sup>96</sup> While best practices suggested that

88. See 5 U.S.C. § 553 (2018).

89. See 49 C.F.R. § 5.17 (2022). Rules are significant if, for example, they have an annual effect of over \$100 million or create serious inconsistencies with another agency’s priorities. Heightened requirements for significant rules were set forth in E.O. 12,866 during the Clinton Administration. See Exec. Order No. 12,866, 3 C.F.R. § 638 (1994).

90. See 49 C.F.R. § 5.17 (2022).

91. *Id.*

92. See Exec. Order No. 12,866, *supra* note 89; see generally Lisa Heinzerling, *Statutory Interpretation in the Era of OIRA*, 33 FORDHAM URB. L.J. 1097, 1098-99 (2006) (giving an overview of the order).

93. See Exec. Order No. 12,866, *supra* note 89 (“‘Agency,’ unless otherwise indicated, means any authority of the United States that is an ‘agency’ under 44 U.S.C. 3502(1), other than those considered to be independent regulatory agencies.”).

94. See Exec. Order No. 13,579, 3 C.F.R. § 13,579 (2012) (“Wise regulatory decisions depend on public participation and on careful analysis of the likely consequences of regulation. Such decisions are informed and improved by allowing interested members of the public to have a meaningful opportunity to participate in rulemaking. To the extent permitted by law, such decisions should be made only after consideration of their costs and benefits (both quantitative and qualitative).”).

95. See OFFICE OF INSPECTOR GENERAL, COST BENEFIT ANALYSIS PROCESS FOR RULEMAKING 5 (2020), [https://www.fdicog.gov/sites/default/files/document/2022-08/20-003eval\\_0.pdf](https://www.fdicog.gov/sites/default/files/document/2022-08/20-003eval_0.pdf).

96. *Id.* at 11.

agencies should conduct varying levels of cost-benefit analysis depending on the significance of the rule at issue, the FDIC was failing to categorize rules.<sup>97</sup> Perhaps due to the lack of a centralized policy, the FDIC had issued some significant rules with no cost-benefit analysis, yet had performed cost-benefit analyses on other, less significant ones.<sup>98</sup>

An algorithm could help address this problem. Imagine an algorithm called the Heightened Procedure Tool (HPT). HPT could be rules-based, meaning it would encode logic like that of the DOT's policies for identifying economically significant rules. Alternatively, if HPT relied on machine learning, the agency could create training data based on previous rules, hypothetical rules tagged by agency employees, or some combination of the two. Of course, the FDIC would need to curate the dataset carefully, lest it encode incorrect logic for categorizing the rules based on previous decisions it believes to be incorrect. HPT might output the same three categories that control cost-benefit analysis under Executive Order No. 12855, or it could make more fine-grained categorizations that allow for a wider degree of variation in the procedures used for a particular rule. While HPT may seem far-fetched, it is not that different from the QDD system described above, which designated certain disability claims for fast-track processing. The challenge for the FDIC would be to create a dataset large enough, given the small number of rulemakings the agency does each year.

What about algorithms acting *as* rules? These are rare. The QDD algorithm was part of a legislative rule, but the details of the algorithm were left vague during notice and comment.<sup>99</sup> Expert agencies already reliant on scientific models might be some of the first to embrace algorithms as part of rules. Cary Coglianese and David Lehr describe one possibility, hypothesizing that the National Machine Fisheries Service might replace the existing statistical models it uses to predict species extinction when formulating new rules with machine learning algorithms that do the same.<sup>100</sup> We could take this example one step further and predict that expert scientific agencies might promulgate algorithms like this as rules in and of themselves. Rather than undergo a new rulemaking for each individual species, the agency might promulgate a rule stating that any species predicted to be endangered by the algorithm—or the algorithm and some combination of other factors—will be presumed endangered.

---

97. *Id.* at 12 (“[T]he process lacks a written policy and procedures that instruct and guide the Working Group in employing their professional judgment in determining when and how to perform cost benefit analyses.”).

98. *Id.* at 15-16.

99. See Engstrom & Ho, *supra* note 56, at 838.

100. Coglianese & Lehr, *supra* note 39, at 45-47.

Finally, agencies also use algorithms in enforcement. At the federal level, the Securities and Exchange Commission (SEC) relies on algorithms to help identify potential insider trading.<sup>101</sup> One tool, called the Advanced Relational Trading Enforcement Metrics Investigation System (ARTEMIS), uses natural language processing, a machine learning technique, to identify changes that might warrant investigation.<sup>102</sup> The SEC also employs algorithms to detect fraud in annual financial forms.<sup>103</sup> A supervised machine learning algorithm classifies investment advisors as high, medium, or low priority for investigation based on a dataset of past registrants who had been referred to the agency's enforcement division.<sup>104</sup> More locally, New York City and Boston city agencies rely on predictive algorithms to allocate inspectors to restaurants to check that they are complying with fire safety and restaurant sanitation codes.<sup>105</sup>

Even though enforcement decisions are not final agency actions, algorithms used in enforcement can be of tremendous consequence. A recent study of algorithms used by the Internal Revenue Service (IRS) to identify candidates for tax audits found that these algorithms contributed to the IRS auditing Black taxpayers at 2.9 to 4.7 times the rate of non-Black taxpayers.<sup>106</sup> This speaks to a broader point about agencies' use of algorithms—these algorithms are not inconsequential just because agencies rarely involve them in final agency action. At present, agencies are far from treating algorithms like all-knowing judges, notwithstanding popular accounts suggesting the possibility.<sup>107</sup> Nevertheless, algorithms already have system-wide effects.

### *C. Agencies' Use of Policies*

Policies structure all kinds of agency action. Typically, policies instruct decision makers on what factors to consider, but they can also afford them full discretion over parts of a decision. One example of a policy is a legislative rule. Legislative rules, which must undergo notice and

---

101. See GOVERNMENT BY ALGORITHM, *supra* note 11, at 23.

102. *Id.*

103. *Id.*

104. *Id.*

105. See Susan Athey, *Beyond Prediction: Using Big Data for Policy Problems*, 355 SCI. 483, 484 (2017).

106. See Hadi Elzayn, Evelyn Smith, Thomas Hertz, Arun Ramesh, Robin Fisher, Daniel E. Ho, & Jacob Goldin, *Measuring and Mitigating Racial Disparities in Tax Audits 3* (Stan. Inst. for Econ. Pol'y Rsch. 3, Working Paper, 2023), <https://siepr.stanford.edu/publications/measuring-and-mitigating-racial-disparities-tax-audits>.

107. See Alexander Stremitzer, Benjamin M. Chen, & Kevin Tobia, *ChatGPT and the Law: Would Humans Trust an A.I. Judge? Yes.*, SLATE (Feb. 28, 2023, 9:40 AM), <https://slate.com/news-and-politics/2023/02/chatgpt-law-humans-trust-ai-judges.html>.

comment, can impose binding obligations on, and create rights for, members of the public. They are critical in adjudications, which often turn on whether regulated parties have violated a legislative rule. More common than legislative rules, however, are guidance documents. Guidance documents vastly outnumber other kinds of policies and can take many forms.<sup>108</sup> Although such guidance cannot be applied inflexibly, it is still useful insofar as it puts regulated parties on notice about an agency's priorities and understanding of its own regulations. Policies such as guidance documents are critical not only to regulated parties, but also to internal agency management. They allow agency staff to supervise lower-level employees, direct discretion from above, and promote uniform application of regulatory and statutory requirements.<sup>109</sup> In this sense, they are the lifeblood of what other scholars have called "internal administrative law."<sup>110</sup> That category encompasses "measures generated by agencies to control their own actions and operations" that are "aimed primarily at agency personnel."<sup>111</sup>

#### *D. Algorithms as Policies*

Algorithms are like policies. To show as much, this Part provides intuition for the argument by discussing one of the most famous—and algorithm-like—policies in administrative law: the SSA's medical-vocational guidelines for disability adjudication. This Part will explain this policy and then consider how a rules-based or machine learning algorithm could be used for a similar purpose. The discussion will then become more abstract and consider commonalities between algorithms and policies. Then, this Part explains that, in practice, algorithms often take the place of existing policies. Other algorithms are trained on datasets based on the consistent application of policies over several years, while still others are helping with tasks where policies previously served important roles. This Part will conclude by explaining the implications of the algorithms-as-policies view for this Article's central focus, reason-giving.

---

108. See Nicholas R. Parrillo, *Federal Agency Guidance and the Power to Bind: An Empirical Study of Agencies and Industries*, 36 YALE J. REG. 165, 168 (2019) ("Nobody knows exactly how much guidance there is, because it is not comprehensively collected anywhere, but its page count for any given agency is estimated to dwarf that of actual regulations by a factor of twenty, forty, or even two hundred.").

109. See Alexander Nabavi-Noori, *Agency Control and Internally Binding Norms*, 131 YALE L.J. 1278, 1287 (2022).

110. See, e.g., Gillian Metzger & Kevin Stack, *Internal Administrative Law*, 115 MICH. L. REV. 1239 (2017) (offering a conceptual and historical account of internal administrative law); MASHAW, *supra* note 52, at 15 (coining the term and calling for scholars to make internal administrative law a research focus). See generally NICHOLAS R. PARRILLO, *ADMINISTRATIVE LAW FROM THE INSIDE OUT* (Nicholas R. Parrillo ed., 2017) (collecting essays on the subject).

111. Metzger & Stack, *supra* note 110, at 1254.

## 1. Familiarizing Algorithms: The Medical-Vocational Grid

The medical-vocational guidelines are one of the most famous policies of all time. These guidelines, often known as the “grid” because of how they are displayed (see Figure 1 below), standardized the process by which agencies determined whether claimants were eligible for disability benefits.<sup>112</sup> Frontline adjudicators determined a claimant’s age, education, work experience, and residual functional capacity. Based on these factors, the grid directed a determination of disabled or not disabled.<sup>113</sup>

**TABLE NO. 2—RESIDUAL FUNCTIONAL CAPACITY:  
MAXIMUM SUSTAINED WORK CAPABILITY  
LIMITED TO LIGHT WORK AS A RESULT  
OF SEVERE MEDICALLY  
DETERMINABLE IMPAIRMENT(S)<sup>a</sup>**

| Rule       | Age                                  | Education  | Previous work experience                            | Decision      |
|------------|--------------------------------------|--|---|---------------|
| 202.01 ... | Advanced age ....                    | Limited or less .....  | Unskilled or none ....                              | Disabled.     |
| 202.02 ... | ..do .....                           | ..do .....   | Skilled or semiskilled—<br>skills not transferable. | Do.           |
| 202.03 ... | ..do .....                           | ..do .....   | Skilled or semiskilled—<br>skills transferable.     | Not disabled. |
| 202.04 ... | ..do .....                           | High school graduate or<br>more—does not provide<br>for direct entry into<br>skilled work. | Unskilled or none ....                              | Disabled.     |
| 202.05 ... | Advanced age ....                    | High school graduate or<br>more—provides for<br>direct entry into skilled<br>work.         | Unskilled or none ....                              | Not disabled. |
| 202.06 ... | ..do .....                           | High school graduate or<br>more—does not provide<br>for direct entry into<br>skilled work. | Skilled or semiskilled—<br>skills not transferable. | Disabled.     |
| 202.07 ... | ..do .....                           | ..do .....   | Skilled or semiskilled—<br>skills transferable.     | Not disabled. |
| 202.08 ... | ..do .....                           | High school graduate or<br>more—provides for<br>direct entry into skilled<br>work.         | Skilled or semiskilled—<br>skills not transferable. | Do.           |
| 202.09 ... | Closely approaching<br>advanced age. | Illiterate or unable to<br>communicate in English.   | Unskilled or none ....                              | Disabled.     |
| 202.10 ... | ..do .....                           | Limited or less—At least<br>literate and able to<br>communicate in English.                | ..do .....  | Not disabled. |
| 202.11 ... | ..do .....                           | Limited or less .....  | Skilled or semiskilled—<br>skills not transferable. | Do.           |
| 202.12 ... | ..do .....                           | ..do .....   | Skilled or semiskilled—<br>.....                    | Do.           |

Figure 1: Residual Functional Capacity.<sup>114</sup>

112. See John J. Capowski, *Accuracy and Consistency in Categorical Decision-Making: A Study of Social Security’s Medical-Vocation Guidelines—Two Birds with One Stone or Pigeon-Holing Claimants?*, 42 MD. L. REV. 329, 384 (1983).

113. Rules for Adjudicating Disability Claims in Which Vocational Factors Must Be Considered, 43 Fed. Reg. 9284, 9285 (proposed Mar. 7, 1978) (codified at 20 C.F.R. pts. 404, 416), <https://www.govinfo.gov/content/pkg/FR-1978-03-07/pdf/FR-1978-03-07.pdf> [hereinafter Rules for Adjudicating Disability Claims].

114. See Capowski, *supra* note 112, at 384.

The SSA promulgated the grid as a legislative rule in 1978.<sup>115</sup> When the SSA created the grid, the disability determination process was the largest system of administrative adjudication in the western world—and it still is.<sup>116</sup> The grid consolidated several fragmented policies into one.<sup>117</sup> Although there were previously policies in place, decisions were plagued by inconsistency in the old system. A study by the National Center for Administration Justice noted that “it is widely believed that the outcome of cases depends more on who decides the case than what the facts are.”<sup>118</sup> In 1977, the *New York Times* reported that the Social Security Disability Insurance (SSDI) program had become “the most arbitrary of the Government’s programs to help the needy, one in which poor people in similar circumstances often receive vastly different treatment.”<sup>119</sup> The grid emerged out of a study by Jerry Mashaw on the SSA disability adjudication process, in which Mashaw advocated for greater use of systemic policies in light of the problems with the existing system.<sup>120</sup> As Jon C. Dubin writes, the grid epitomizes a form of what Mashaw has called “bureaucratic rationality,” which seeks to promote efficiency, consistency, and accuracy through centrally formulated policies, possibly at the expense of greater individualized evaluation.<sup>121</sup>

Although the existing system had problems, the grid was also controversial. During the rulemaking process, commenters worried that the grid discriminated against young people,<sup>122</sup> did not account for whether jobs were available in the claimant’s place of residence,<sup>123</sup> and did not accurately measure the claimant’s level of education.<sup>124</sup> Others objected to grid schemes on more general, procedural grounds. They worried that human judgment was giving way to bureaucratic procedure.

---

115. *Id.* at 9291.

116. See Nicholas R. Parrillo, *Jerry L. Mashaw's Creative Tension with the Field of Administrative Law*, in *ADMINISTRATIVE LAW FROM THE INSIDE OUT* 1, 2 (Nicholas R. Parrillo ed., 2017).

117. See Rules for Adjudicating Disability Claims, *supra* note 113, 9294 (1978) (“As now explained elsewhere in this preamble, the proposed regulations do, in fact, represent a consolidation and elaboration of longstanding medical-vocational evaluation policies which have their bases in the Social Security Act and are in accordance with Congressional intent. These policies, however, have heretofore been reflected in fragmented guides and have not been readily available in the same format at all levels of adjudication.”).

118. Capowski, *supra* note 112, at 343.

119. David E. Rosenbaum, *Huge Federal Disability Program Faces Inequities, Fund Woes, Suits*, N.Y. TIMES, July 27, 1977, <https://www.nytimes.com/1977/07/27/archives/huge-federal-disability-program-faces-inequities-fund-woes-suits.html>.

120. See Jon C. Dubin, *Overcoming Gridlock: Campbell After a Quarter-Century and Bureaucratically Rational Gap-Filling in Mass Justice Adjudication in the Social Security Administration’s Disability Programs*, 62 ADMIN. L. REV. 937, 940 (2010).

121. *Id.*

122. See Rules for Adjudicating Disability Claims, *supra* note 113, at 9291.

123. *Id.* at 9295.

124. *Id.* at 9299.



Some believed the agency was resorting to an “average man” concept to make decisions rather than assessing claimants case-by-case.<sup>125</sup> What effect would the grid have, they asked, on under-resourced claimants, who might lack the knowledge or authority to protest attempts to “force them into a ‘pigeon-hole’”?<sup>126</sup> Would claimants need attorneys to help them navigate the grid to make viable claims for disability benefits?<sup>127</sup> Such fears of “computerized adjudication” abounded.<sup>128</sup>

The point in mentioning this decades-old controversy is not to make commenters’ concerns seem naïve or antiquated. Rather, it is to make the problems posed by algorithms feel more familiar. Agencies can and do use policies in place of fully individualized consideration, both for internal and external decisions—but doing so has costs. Rather than relitigate this longstanding debate between systemic control and individualized decision making, this Article’s goals are more modest. At least in some cases, tools of bureaucratic management are necessary. The question is whether these tools necessarily must take the form of policies instead of algorithms.

Setting aside the reason-giving issue for a moment, the grid can help shore up intuition about the similarities between policies and algorithms. After all, the grid is an algorithm itself. An algorithm takes a set of inputs and, based on a finite sequence of instructions, produces an output. Indeed, if the SSA chose, it could convert the grid into a rules-based algorithm (see Figure 2). Thus, in many cases, the only thing separating a rules-based algorithm and a policy like the grid is digitization.

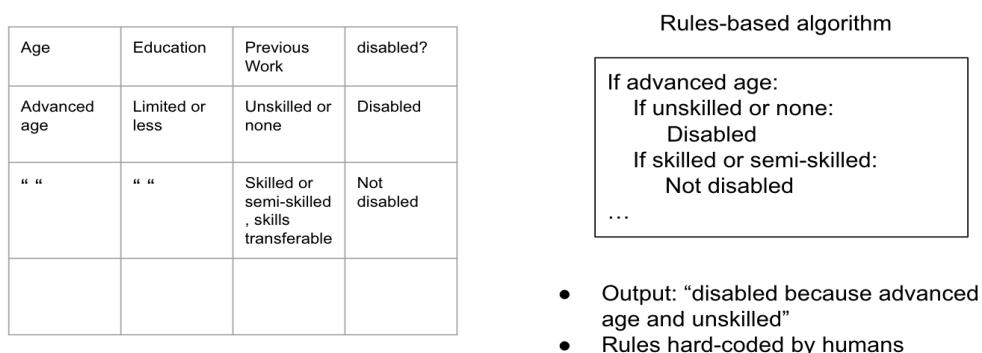


Figure 2: Converting the medical-vocational grid (left) into a rules-based algorithm (right).

125. *Id.* at 9291.

126. *Id.*

127. *Id.* at 9301.

128. *Id.* at 9291.

Granted, the algorithm-as-policy argument is trickier to grasp for machine learning algorithms because the underlying rules of the algorithm are not readily visible or intuitive. As a start, though, consider a hypothetical machine learning algorithm that would replace the medical-vocational grid. If such an algorithm feels far-fetched, remember that the QDD algorithm discussed above essentially does this, except instead of rendering a final agency decision, it fast-tracks some claimants for faster processing.<sup>129</sup> This hypothetical algorithm would be trained on past disability adjudications involving the grid. It could consider the same four features as the grid but might include others as well. Additionally, it might consider some features more granularly. So, for example, it might account for a claimant's exact age rather than slotting them into a category such as "advanced age." The basic takeaway is that such a machine learning algorithm could assume the same position within the SSA as the grid. Functionally, they would be identical. With that intuition established, this Part considers more precisely what algorithms and policies have in common.

## 2. Breaking it Down: What Algorithms and Policies Have in Common

At a higher level, what exactly do algorithms and policies have in common? Algorithms and policies help agencies achieve similar goals (consistency, accuracy, and efficiency), require similar human labor on the front end, and have similar effects within agencies when deployed.

First, goals. Agencies typically turn to algorithms and policies to make decisions more efficient, consistent, and accurate.<sup>130</sup> Algorithms and policies can promote efficiency because they can scale to meet the needs of growing bureaucracies. Hiring more employees to do an existing job does not require creating a new algorithm any more than it requires drafting a new analog policy. Once an algorithm or policy is in place, it can guide the decisions of new employees performing the same task. Algorithms and policies can reduce the amount of effort individual decision makers must expend to come to a decision. With algorithms or

---

129. *See supra* Section I.B.

130. In his report for ACUS, Coglianese listed accuracy, capacity, speed, and consistency as possible reasons to use algorithms. *See* CARY COGLIANESE, A FRAMEWORK FOR GOVERNMENTAL USE OF MACHINE LEARNING 34-37 (2020). On the policy side, Jerry Mashaw has been a vocal advocate for the pursuit of these goals. *See, e.g.*, MASHAW, *supra* note 52, at 26 ("And, of course, this application of knowledge must in any large-scale program be structured through the usual bureaucratic routines: selection and training of personnel, detailed specification of administrative tasks, specialization and division of labor, coordination via rules and hierarchical lines of authority, and hierarchical review of the accuracy and efficiency of decisionmaking.").

policies in place, agency employees do not need to reinvent the wheel to solve every problem. Even if they are not binding on agency employees, algorithms and analog policies give agency employees a place to start.

As for consistency, agencies have long justified centralized policies by pointing to disparities in outcomes when decisions are left to individual discretion. The SSA's notice of proposed rulemaking for the medical-vocational grid is a case in point. As the agency wrote, "It is expected that publication of the proposed amendments will result in a higher degree of consistency and equity of medical-vocational adjudication throughout the country at all adjudicative levels."<sup>131</sup> A desire for more consistent outcomes is also one potential reason for turning to algorithms and policies. None of this is to say algorithms and policies are always accurate, consistent, or efficient; it is only to say that those values are a driving force for agencies that turn to them.

Next, labor. Agencies must invest significant employee time, research, and money on the front end to create quality algorithms and policies. Not closely resembling an individual person's knowledge, algorithms and policies instead reflect the wisdom of the bureaucratic organization. Creating algorithms and policies often requires tremendous amounts of research and information—more than an individual person can meaningfully absorb. For example, the SSA built its algorithms for disability adjudication based on decades of data, research, and prior policies.<sup>132</sup> Similarly, the SSA's medical-vocational grid consolidated disparate policies that had previously governed disability adjudication.<sup>133</sup> Machine learning algorithms are particularly labor intensive, as they require large datasets that are created and cleaned by large teams of people.

Beyond requiring many people's labor, algorithms and policies are similar in that they encode the work of domain experts. Enforcement policies at the Environmental Protection Agency (EPA), for instance, give frontline employees precise cutoffs for contaminants based on research by scientists who do not themselves carry out enforcements. This, too, is one potential benefit of algorithms: because technical employees often—and one may add, should—build them in consort with domain experts,<sup>134</sup> algorithms can diffuse domain expertise throughout an organization.

Finally, effects. It helps to start with what algorithms and policies

131. See Rules for Adjudicating Disability Claims, *supra* note 113, at 9297.

132. See Kurt Glaze, Daniel E. Ho, Gerald K. Ray, & Christine Tsang, *Artificial Intelligence for Adjudication: The Social Security Administration and AI Governance*, in OXFORD HANDBOOK OF AI GOVERNANCE (Justin B. Bullock ed., 2022), <https://academic.oup.com/edited-volume/41989/chapter/355439450> [hereinafter *The Social Security Administration and AI Governance*].

133. See Capowski, *supra* note 112, at 346 n.90.

134. See Glaze et al., *supra* note 132.

usually do *not* do, which is remove humans from the decision making process entirely. Instead, they operate in conjunction with them. Both are built by humans. As others have written, “[e]ven when we ‘replace’ a human with an algorithm for the purpose of reaching an individual decision, humans designed the system, asked the question, and often implemented the conclusion. It’s humans all the way down.”<sup>135</sup> Not only is this true at the design stage, but it often remains true throughout the algorithm’s lifetime. But unlike humans, neither algorithms nor policies “think” in any familiar sense of the word. Even when algorithms exert significant influence over human decisions, there is often a human in the loop who renders a final decision or at least monitors the output of the algorithm.<sup>136</sup> That mirrors the status of policies within agencies. While policies are sometimes binding on agency employees—in which case they may need to undergo notice and comment—other times employees must formulate independent reasons for a decision even when a policy is in place.

Both algorithms and policies influence the decisions of many people within an agency. In a fictional world of full discretion, agency employees would make decisions on an individual basis based on criteria of their choosing. But through policies and algorithms, bureaucratic managers shape decisions from above. And because policies and algorithms centralize decision making, they can create focal points for review. These reviewers might be federal court judges, agency managers, or members of the public. For instance, creating the medical-vocational grid involved a notice and comment process, during which members of the public could challenge the grid for potentially discriminating against young people,<sup>137</sup> failing to account for a claimant’s place of residence,<sup>138</sup> and inaccurately measuring a claimant’s level of education.<sup>139</sup> Similarly, because algorithms can be audited, it is possible to surface potentially biased or erroneous decision making at the system level. For as much as algorithms and policies undermine reason-giving, they can also facilitate more generalized contestation of agency decisions.

Ultimately, algorithms and policies are alike in their goals, in the labor required to build them, and in their effects within agencies. It is no surprise, then, that agencies often turn to algorithms where they previously used policies.<sup>140</sup> Today, the SSA leads federal agencies in the

---

135. Crootof et al., *supra* note 49, at 443.

136. *Id.*

137. *See* Rules for Adjudicating Disability Claims, *supra* note 113, at 9291.

138. *Id.* at 9295.

139. *Id.* at 9299.

140. *See supra* Section I.D.2.

adoption of algorithms.<sup>141</sup> One reason is that the SSA had already invested heavily in creating structured policies and procedures for decades by the time machine learning technology became available.<sup>142</sup> These included the grid, of course, but also more internally facing policies, such as questionnaires that guided adjudicators through adjudications.<sup>143</sup> In addition, the SSA collected and organized data related to the outcomes of disability adjudications to improve the quality and consistency of these determinations.<sup>144</sup> These policies and datasets formed a critical “foundational infrastructure” for creating algorithms.<sup>145</sup> That is unsurprising once we realize algorithms are policies of a different kind.

### 3. Two Implications for Reason-Giving

Agencies adopt policies because they are necessary for bureaucratic management. While one can and should debate the merits of a particular policy, agencies cannot forgo policies altogether. Policies are the difference between ad hoc decision making and the work of administration. Once we recognize the necessity of at least *some* policies, envisioning a place for algorithms within an agency becomes easier because algorithms serve a similar end: bureaucratic administration.

The algorithms-as-policies view has two important implications for a normative and legal assessment of algorithmic reason-giving. First, it shifts the normative baseline. Given that algorithms’ closest counterparts are policies, not people, comparing algorithmic reason-giving to what we might call policy reason-giving is the best way to tease out what is distinctive about algorithmic reason-giving relative to the status quo. It is unnecessary to consider the much more difficult comparison between algorithms and humans to answer the reason-giving question. So long as one accepts that policies have some role to play in the administrative state—debating the proper scope of this role is an entire research agenda of its own—one can ask a simpler question: do the issues algorithms pose for reason-giving mean that they should not, whether normatively or legally, be able to play the same role as policies?

The second implication of the algorithms-as-policies view is that it gives us a framework for evaluating algorithmic reason-giving. Policies show that reason-giving entails two different things: first, the giving of reasons *for* a policy or algorithm; and second, the giving of reasons *based on* a policy or algorithm. The next Section assesses differences between

---

141. See Glaze et al., *supra* note 132.

142. *Id.*

143. *Id.*

144. *Id.*

145. *Id.*

policies and algorithms with respect to the first kind of reason-giving. Section III takes on the second kind.

## II. SYSTEMIC REASONS FOR ALGORITHMS/POLICIES

When creating algorithms and policies, agencies must explain why they make certain choices. In giving these explanations, agencies offer systemic reasons for the algorithm or policy in question. Systemic reasons, whether for algorithms or policies, can lead to higher quality decision making, make agencies more accountable for their decisions, and protect the dignitary interests of regulated parties. The greatest downside of systemic reasons for algorithms is that many of them are technically complex, which makes them harder to scrutinize. However, technically complex systemic reasons are nothing new. On the contrary, many agencies must rely on specialized expertise to give reasons for their decisions. Algorithms are thus most comparable to technically or scientifically complex policies. The challenge of algorithms is that they force agencies to give technically complex systemic reasons where they might not have been required to do so before. This shift has accountability and dignitary costs and may undermine the quality-promoting effect of systemic reasons. Nevertheless, other considerations may mean an algorithm is still the best choice. At most, the inherent technical complexity of systemic reasons for algorithms nudges the balance in favor of policies where all other factors are equal.

### *A. What Do Systemic Reasons Look Like?*

Fundamentally, a reason is an answer to the question, “why?” Reasons may answer a variety of different “why” questions. Why regulate this area? Why adopt this policy and not that one? Why these technical methods and not competing ones? For any given “why,” there are numerous reasons an agency could give. Reasons may rely on technical findings, or an agency might offer reasons grounded in politics. This Part puts the discussion of systemic reason-giving on a more concrete footing by discussing various reasons an agency might give for using an algorithm or policy. It primarily relies on two examples—one an algorithm, the other a policy—to explore the nature of systemic reasons. Although some systemic reasons for algorithms are distinctive, such as reasons given for choosing a particular machine learning model, many kinds of systemic reasons for algorithms have a clear analogue in systemic reasons for policies. When comparing systemic reasons for algorithms to those for policies, then, the number of reasons we should have concerns about is smaller than one might expect.

First, consider the systemic reasons the Allegheny County Department of Human Services (DHS) gave when designing the Allegheny Family Screening Tool (AFST), a tool built to improve child-welfare screening decisions. Although the DHS is a state agency, the AFST is a useful example because the DHS prepared not one, but two reports discussing the choices it made when designing the AFST. The final report included a description of how the DHS developed the tool, an ethical analysis of the use of predictive models for this purpose, an impact evaluation summary, and responses to frequently asked questions.<sup>146</sup> This report was effectively a 241-page catalogue of systemic reasons for the algorithm.

One category of systemic reasons for algorithms is those given for the technical choices an agency makes when designing an algorithm. For example, the DHS gave reasons for choosing a particular machine learning model,<sup>147</sup> for training the model on certain features,<sup>148</sup> and for including data in the training and test datasets.<sup>149</sup> Often, the DHS expressed these reasons in terms that non-technical audiences could understand. Consider the agency's decision to use a Least Absolute and Shrinkage Operator (LASSO) regression model. One needs machine learning knowledge to understand what it means when the DHS said the LASSO model "was trained on the training partition using [ten]-fold cross-validation."<sup>150</sup> Yet, that does not mean the choice was completely unexplainable to those without machine learning expertise. Many of the agency's reasons for choosing the LASSO model make intuitive sense even without technical knowledge. For instance, the DHS justified choosing the LASSO model because it had the best combination of predictive success, parity between racial groups, and ease of implementation.<sup>151</sup> Those reasons track important values in

---

146. Rhema Vaithianathan, Emily Putnam-Hornstein, Nan Jiang, Parma Nand & Tim Maloney, *Developing Predictive Risk Models to Support Child Maltreatment Hotline Screening Decisions: Allegheny County Methodology and Implementation*, in DEVELOPING PREDICTIVE RISK MODELS TO SUPPORT CHILD MALTREATMENT HOTLINE SCREENING DECISIONS, 14 (2019) [hereinafter AFST REPORT].

147. *Id.* ("[Decision-tree] methods have the advantage that they are often more accurate – with higher precision, recall and area under the ROC.").

148. *Id.* at 15 ("After an independent ethical review of this project and lengthy discussions between community stakeholders, internal staff, and members of the research team, the County made the decision that race could be included as a predictor variable if it substantively improved the predictive accuracy of the model.").

149. *Id.* at 11 n.5 ("The cut-off date [for data in the dataset] was determined by the fact that Allegheny County transitioned to its current KIDS data system in 2008.").

150. Rhema Vaithianathan et al., *Allegheny Family Screening Tool: Methodology, Version 2*, in DEVELOPING PREDICTIVE RISK MODELS TO SUPPORT CHILD MALTREATMENT HOTLINE SCREENING DECISIONS, 10 (2019).

151. *Allegheny Cnty. Dep't of Hum. Servs., Frequently-Asked Questions*, in DEVELOPING PREDICTIVE RISK MODELS TO SUPPORT CHILD MALTREATMENT HOTLINE SCREENING DECISIONS, 19-20 (2019). ("To determine which methodology would be used, researchers considered 1) overall performance

administration: accuracy, consistency, and efficiency, respectively. One needs little specialized expertise to understand them, even though the DHS used them to justify technical decisions.

Before technical considerations are even an issue, agencies give reasons for addressing a particular problem and pursuing certain policy goals. Agencies then judge an algorithm's success based on how much it helps the agency realize these goals. In the case of the AFST, the DHS explained that it wanted to make decisions more efficiently and consistently, put county resources towards the most vulnerable populations, and improve the health and safety of county residents.<sup>152</sup> With those considerations in mind, the DHS chose to build a tool to aid welfare screeners. Specifically, it built an algorithm to predict the likelihood of longer-term adverse events, as opposed to predicting the likelihood of current abuse and neglect based on screening calls.<sup>153</sup> The reason the DHS gave for this goal was that, because screening staff focused on information immediately in the referral, an algorithm predicting longer-term outcomes would complement, rather than repeat, the frontline workers' process.<sup>154</sup> This decision was debatable. One may think the algorithm should be redundant with the human decision maker, to ensure that human and algorithmic decisions align. Regardless, DHS gave reasons for its choice of goal.<sup>155</sup> These reasons were not specific to algorithms, but were grounded in the agency's broader needs and priorities.

Ongoing auditing can make for some of the best systemic reasons. The first version of the AFST operated from 2016 to 2018, over which time the DHS studied its results.<sup>156</sup> Because Version One failed to help the agency reach some of its intended goals, the DHS released Version Two based on a study of decisions made in conjunction with the algorithm.<sup>157</sup> The DHS changed both the target outcome of the algorithm,<sup>158</sup> and some

---

and accuracy for the high-risk groups; 2) accuracy for [B]lack children versus non-[B]lack children; 3) ease of implementation and quality checking; and 4) whether the model showed a positive correlation between the score generated and the probability that the child would be involved in a fatality or near-fatality [fifty] days or more after the score was generated.”).

152. AFST REPORT, *supra* note 146, at 4.

153. *Id.* at 8.

154. *Id.*

155. *Id.* at 4.

156. *Id.* at 18.

157. Version 1 did not lead to increases in the rate of children screened for investigation, did not lead to decreases in re-referral rates for children screened-out without an investigation, and did not clearly result in greater screening consistency between screening callers. However, Version 1 did increase the identification of children in need of child welfare intervention and reduced case opening disparities between white and Black children. Allegheny Cnty. Dep't of Hum. Servs., *supra* note 151, at 17-18.

158. *Id.* at 19 (“AFST Version 1 (V1) was designed to predict: 1) the likelihood a child would experience abuse or neglect serious enough to be placed in an out-of-home setting within two years of the



of the data sources used in the algorithm.<sup>159</sup> It also switched from a primitive regression model to a machine learning model. The agency justified these changes in depth.<sup>160</sup>

Given the range of forms policies can take, generalizing about systemic reasons for policies is more challenging. Consider the reasons the SSA gave in a notice of proposed rulemaking when designing the medical-vocational grid.<sup>161</sup> First, the agency gave reasons for using a grid in the first place.<sup>162</sup> The SSA reasoned that the grid would: make clearer to claimants and their lawyers how the agency determined disability, make determinations more consistent and accurate, and promote better understanding and acceptance by the public and the courts of disability determinations.<sup>163</sup> In other words, the reasons for using the grid were the benefits of using any policy: consistency, accuracy, and transparency.

Having established why it was relying on a grid, the SSA gave reasons for its more specific choices. Like the DHS explaining why it used certain features in an algorithm, the SSA explained why it made disability determinations based on only age, vocation, and work experience. For one, the underlying statute required it.<sup>164</sup> For another, experience with disability adjudication showed the importance of considering these factors.<sup>165</sup> Reasons grounded in research supported relying on these factors, too.<sup>166</sup>

---

initial call if the call were screened-in for investigation and 2) the likelihood there would be a re-referral to the hotline within two years if the call were screened-out. Based on feedback from staff and external validation of the model using hospitalization data, we determined that the scores from the re-referral model were not as strongly related to the key outcome of concern, serious abuse and neglect. AFST Version 2 (V2) therefore only predicts the likelihood of out-of-home placement within two years.”)

159. AFST REPORT, *supra* note 146, at 4 (“Public benefits data were excluded as the current data feeds no longer align to the historic data used to develop V1. Some behavioral health records were eliminated because of temporal variability. In addition, variables regarding the current allegations on the referral were added at the request of call-screening staff.”)

160. *Id.* at 3.

161. *See* Rules for Adjudicating Disability Claims, *supra* note 113, at 9284.

162. *Id.* (“In publishing the proposed amendments, the Social Security Administration intends to consolidate and elaborate upon longstanding medical-vocational evaluation policies for adjudicating disability claims in which an individual’s age, education, and work experience must be considered in addition to the medical condition.”).

163. *Id.* at 9285.

164. *Id.* (“[B]ecause of the clearly limited statutory definition, those factors which relate primarily not to disability but to an individual’s ability to obtain employment have been excluded from consideration.”).

165. *Id.* at 9289 (“Prior experience of the Social Security Administration in determining when age makes a difference in disability determinations has also been considered . . .”).

166. *Id.* at 9290 (“Reasoning ability would affect the individual’s ability to follow instructions and make judgments in a work situation. Language competence relates to ability to read, write, and speak. The inability to meet the language requirements at an elementary level would restrict even the number of unskilled jobs a person would be able to do. Similarly, the inability to perform single calculations in addition and subtraction would represent vocational restrictions in performing some unskilled jobs.”).

The SSA then responded to comments from the public. In response to a charge that the agency was resorting to “computerized adjudication,” the SSA gave a simple reason for relying on centralized standards—namely, that it would be impossible to administer the social security program without them.<sup>167</sup> The SSA also rebutted comments saying that the agency lacked the authority to promulgate regulations like the medical-vocational grid.<sup>168</sup> Having addressed broader concerns about the program, the SSA responded to more granular comments—why the agency relied on national job statistics rather than local ones (the statute required it),<sup>169</sup> why the guidelines considered an individual’s existing job skills (skills make it easier to find a new job),<sup>170</sup> why the agency considered the last fifteen years of an individual’s work experience (jobs were presumed to change in about that time period),<sup>171</sup> and why individuals over forty were still considered “young” (data and prior experience suggested that major problems do not appear until around fifty).<sup>172</sup>

Comparing the reasons given for the AFST and for the medical-vocational grid, similarities emerge. In both cases, agencies gave reasons for many similar “why” questions: why make decisions based on some factors but not others? Why rely on some data but not others? And why create the algorithm or policy in the first place? In addition, the reasons the agencies gave for their choices relied on a combination of secondary research, agency experience, and pragmatic considerations.

That said, systemic reasons for algorithms are distinctive in that they are sometimes unavoidably technical. While technical reasons are, of course, not foreign to reason-giving for policies, relying on algorithms will inject a need for technical expertise into less technical areas of administration. Nonetheless, many, if not most, of the systemic reasons given for the AFST were not grounded in machine learning minutiae. Even technical decisions were often justified in domain-specific—here, the domain being child welfare—terms, the very terms an agency would use to justify a policy relating to child welfare.

In some ways, the AFST is the gold-standard for systemic reason-giving for algorithms. The state DHS published two comprehensive reports and responded to criticism. Not all algorithms are ventilated so.

---

167. *Id.* at 9291.

168. *Id.* at 9293 (“The statute, thus, bestows very broad authority upon the Secretary to issue regulations to discharge his responsibilities for administering the social security programs, subject only to the limitation that the Secretary’s rules, regulations and procedures should not be ‘inconsistent’ with the provisions of the Act.”).

169. *Id.* at 9295.

170. *Id.* at 9298.

171. *Id.* at 9299.

172. *Id.* at 9300.

For example, the SSA publicized reasons for incorporating the QDD algorithm into its workflow, but gave no reasons related to any specific implementation of an algorithm.<sup>173</sup> Judging by volume alone, the DHS's reason-giving for the AFST far surpassed the SSA's reason-giving for medical-vocational guidelines. Perhaps this was because the DHS was aware that introducing an algorithm into the child-welfare process would be controversial.<sup>174</sup>

At a minimum, the AFST illustrates that quality systemic reason-giving for algorithms is possible. Moreover, the comparison with the medical-vocational guidelines suggests that algorithmic reason-giving may not differ significantly from systemic reason-giving for policies. The next two Parts argue that any differences are not so important as to justify rejecting algorithms wholesale.

### *B. Virtues of Systemic Reasons for Policies*

This Part focuses on how systemic reasons for policies promote three virtues: increasing the quality of decisions, fostering accountability, and preserving the dignitary interests of regulated parties. Other benefits exist.<sup>175</sup> This Article focuses on these three virtues because they capture differences between systemic reasons for policies and systemic reasons for algorithms.

#### 1. Promoting Quality Decision Making

Systemic reasons for policies improve decisional quality.<sup>176</sup> Just anticipating the need to give reasons can enhance deliberative rigor.<sup>177</sup> Reason-giving can also reveal when bias is driving a decision and might encourage an agency to engage different groups and conduct research about the relevant issues. Reason-giving is particularly effective at improving decisional quality if agencies make reasons widely known,

---

173. See Amendments to the Quick Disability Determination Process, 20 C.F.R. §§ 404, 405, 416 (2007) (“[T]he specific criteria of the predictive model are not prescribed by this rule, and therefore we are making no changes to this rule in response to this comment.”).

174. It has indeed been controversial. See Huq, *supra* note 36, at 1893 (detailing commentators’ concerns about the AFST, including evidence of racial disparities, the risk of “dehumanization,” and the privacy threats resulting from aggregating data to train the algorithm); EUBANKS, *supra* note 15, at 132-54 (describing the implementation of the AFST and articulating concerns).

175. See, e.g., Deeks, *supra* note 30, at 626-34 (offering five virtues of reason-giving: improving decisional quality, promoting government efficiency, constraining decision-makers, strengthen decision-makers’ legitimacy, and fostering accountability).

176. *Id.* at 627-28.

177. See Mary B. DeRosa & Mitt Regan, *Deliberative Constitutionalism in the National Security Setting*, in THE CAMBRIDGE HANDBOOK OF DELIBERATIVE CONSTITUTIONALISM, at 29 (Jeff King et al. eds., 2018) (“[A]nticipating the need for [reason-giving] also can enhance deliberative rigor.”).

such as by exposing them during notice and comment. That is because parties outside the agency can then scrutinize the grounds for a decision and suggest alternatives. But reasons need not be public to promote quality decision making. As Ashley Deeks has argued, reason-giving within an agency encourages deliberation, exposes ideas to critique, increases the number of inputs considered, and encourages agencies to consider perspectives they otherwise might have neglected.<sup>178</sup>

## 2. Fostering Accountability

Systemic reason-giving for policies can make agencies more accountable to a variety of audiences. When an agency gives a reason for a policy choice, other parties—members of Congress, the President, regulated parties, the courts—can evaluate the reason and make sure it is sound. These other parties can make sure the agency is making decisions based on the legitimate grounds for a policy, which might be specified by statute. Perhaps most importantly for accountability is the fact that reasons are the basis for judicial review of agency action. Courts consult the record to see “what major issues of policy were ventilated . . . and why the agency reacted to them as it did.”<sup>179</sup> At least in theory, “[t]he judicial demand for *reasons* has become a legitimate procedural version of an otherwise illegitimate substantive demand for *reasonableness*, as judicially determined.”<sup>180</sup> In this sense, reason-giving promotes a middle-level of judicial review of agency policies, wherein courts are the arbiters not of the policy but the procedures that produced it. “The agency may make policy choices, so long as it explains how its exercise of discretion is connected to its statutory authority and to the technical facts that have been developed through the rulemaking proceeding.”<sup>181</sup> The extent to which this restrained conception of judicial review matches reality is debatable. Regardless, any review of agency policies, by any relevant stakeholder, will include an evaluation of the reasons an agency gave for the policy.

## 3. Preserving Dignitary Interests

Dignity is hard to pin down.<sup>182</sup> It has already been established that

---

178. See Deeks, *supra* note 30, at 667-70.

179. See Jerry L. Mashaw, *Reasoned Administration: The European Union, the United States, and the Project of Democratic Governance*, 76 THE GEO. WASH. L. REV. 99, 110 (2007) (quoting *Auto. Parts & Accessories Ass’n v. Boyd*, 407 F.2d 330, 338 (D.C. Cir. 1968)).

180. *Id.* at 111.

181. *Id.*

182. See Rachel Bayefsky, *Dignity as a Value in Agency Cost-Benefit Analysis*, 123 YALE L.J. 1732, 1739 (2014) (explaining that dignity can include a status of equality, a feature of individuals with

when an agency does not give a reason for a decision, that increases the risk that the decision could be wrong or that the agency is doing something nefarious. But even if the agency is acting within the bounds of its authority and making quality decisions, failing to give reasons harms the regulated party. The concept of dignity captures this broad category of potential harm to an individual person.

Margot Kaminski and Jennifer Urban have emphasized the importance of dignitary concerns such as a respect for individual selfhood.<sup>183</sup> When agencies give reasons for their policies, they convey respect for the people they govern.<sup>184</sup> Giving reasons tells interested parties that the agency cares that the parties understand the decision, even if they disagree with it. Moreover, in the context of notice and comment rulemaking, giving reasons for the policy fosters a conversation between the agency and interested parties. It tells regulated parties that, to some degree, they can shape the development of the laws that affect them. While the agency will not agree to all parties' suggestions, by responding to these concerns the agency can show that it has at least considered the parties' arguments.

Dignity is closely related to the ability to contest agency action. Contestation furthers dignitary interests because it allows interested parties to have agency in the process through which regulations affecting them are made.<sup>185</sup> Systemic reasons for policies facilitate contestation by giving interested parties more specific targets. Rather than argue with the entire policy, interested parties can contest the specific grounds for the policy. In this way, reason-giving makes contestation more tractable.

An important caveat to this whole discussion is that systemic reasons mostly promote dignity to the extent that they are made public. Most policies developed by agencies govern internal decisions. Only agency employees will see the systemic reasons for these policies, even if the internal policies have external effects.

### *C. How Systemic Reasons for Algorithms Measure Up*

Having considered the virtues of systemic reason-giving for policies, this Part considers whether systemic reason-giving for algorithms has the same benefits. The question is one of degree. At least to some extent, systemic reasons for algorithms promote quality, increase accountability,

---

autonomy, or an element of basic humanity, and that dignitary harm can involve, among other things, the loss of reputation, feelings of humiliation, exposure of private details, and diminished status).

183. See Margot E. Kaminski & Jennifer M. Urban, *The Right to Contest AI*, 121 COLUM. L. REV. 1957, 1993 (2021).

184. See Jodi L. Short, *The Political Turn in American Administrative Law: Power, Rationality, and Reasons*, 61 DUKE L.J. 1811, 1822 (2012) (“[T]he act of giving reasons demonstrates respect for the governed subject.”).

185. See Kaminski & Urban, *supra* note 183, at 1993.

and protect dignitary interests. But do they do so *enough*? This Section concludes that they may. There is significant overlap between systemic reasons for policies and for algorithms. In both cases, agencies give reasons for creating the overall program, for setting certain objectives, and for relying on certain forms of evidence. Other similarities are present as well. But because the task is comparative (to understand how algorithms measure up against policies), this Part will primarily focus on the ways in which systemic reason-giving for algorithms is different from systemic reason-giving for policies, for better or for worse.

### 1. Promoting Quality Decisions

Algorithmic quality, like policy quality, is empirical and case specific. Some algorithms are accurate (or consistent or fair), and others are not. To assess the extent to which systemic reasons for algorithms can be expected to promote quality decisions, we must proceed at a somewhat high level of abstraction and ask what attributes these reasons have that might support or undermine decisional quality. In many cases, systemic reasons for algorithms promote quality decisions for the same reasons as systemic reasons for policies. Both force deliberation, encourage agencies to involve different stakeholders, and lead agencies to articulate specific values and goals when making decisions. Both also facilitate public participation. But systemic reasons for algorithms are novel in ways that have implications for the quality-promoting function of systemic reason-giving. Some of these novelties merit concern, while others suggest that systemic reason-giving for algorithms could be even more quality promoting than systemic reason-giving for policies.

One reason to believe systemic reasons for algorithms promote quality is that algorithms make error rates explicit.<sup>186</sup> Error rates could thus facilitate robust notice and comment, as pointing to error rates would help commenters give reasons for adopting or eschewing a particular algorithm. More modestly, error rates could facilitate better reason-giving within an agency, as different internal stakeholders can be aware of the accuracy of the model against a test dataset. Agencies can point to a low error rate as a *prima facie* good reason for choosing a particular algorithm. Conversely, agencies should give strong counter-reasons for choosing an algorithm with a high error rate, meaning a high error rate might promote further investigation and refinement.

---

186. See Deborah Hellman, *Measuring Algorithmic Fairness*, 106 VA. L. REV. 811, 842 (2020) (arguing that a lack of error rate parity between different groups provides a moral reason to investigate the data further and hesitate to use the data); *Oregon DHS Safety at Screening Tool – Development and Execution*, <https://www.oregon.gov/odhs/data/orrai/safety-at-screening-report.pdf> (Nov. 2019), 7 (explaining the agency’s goal of reducing disparities between the error rates shown for different groups).

Explicit error can also help agencies create algorithms that honor the agencies' stated policy priorities. Errors include both false positives and false negatives.<sup>187</sup> False positive and false negative rates bring policy tradeoffs to the fore by emphasizing the effects of action and inaction. For instance, the DHS recognized that, for the AFST, false positives meant the agency might disrupt more families, while false negatives potentially meant some children would not get support the agency believed they needed. Accordingly, the DHS aimed to minimize both the false positive and false negative rate.<sup>188</sup> One could imagine a different scenario in which an agency was less concerned about the implications of agency action than inaction. In that scenario, an agency could give reasons for minimizing the false negative rate, with less regard for the effects on the false positive rate. In either scenario, false positive and negative rates force a reckoning with the likely effects of an agency policy and thus could make algorithms more forward-looking. Granted, not every error can easily be categorized as a false positive or negative. For decisions that are not binary, an algorithm could give a wildly inaccurate estimate of the costs of something, of how long something might take, or of how much assistance a person might need. In these cases, agency employees could examine the magnitude of error, the direction of error (too high or too low), and other metrics to refine the algorithm.

That said, error rate as measured by performance against a test dataset is only part of the picture. The quality of the underlying data matters, too. One issue related to data quality is that the dataset might exclude important data points. Once an algorithm is in the wild, it will encounter unexpected cases. Some of these cases might be hard for algorithms to deal with; scholars refer to this as the "long-tail problem."<sup>189</sup> Margot Kaminski and Jennifer Urban recounted one example at the United States Department of Agriculture (USDA), where a fraud-alert algorithm identified transactions at a Somali-American grocery store as fraudulent solely because customers there bought items in whole-dollar amounts.<sup>190</sup> The long-tail problem exists for policies, too; a request for an exception is essentially an argument that a policy, whatever its virtues, is inapt for a particular situation. Still, robust reason-giving should include at least

---

187. See Huq, *supra* note 36, at 1915-16.

188. See Vaithianathan et al., *supra* note 153, at 10.

189. Evan Ackerman, *Autonomous Vehicles vs. Kangaroos: The Long Furry Tail of Unlikely Events*, IEEE SPECTRUM (July 5, 2017), <https://spectrum.ieee.org/cars-that-think/transportation/self-driving/autonomous-cars-vs-kangaroos-the-long-furry-tail-of-unlikelylevents>.

190. See Kaminski & Urban, *supra* note 183, at 1973 (citing H. Claire Brown, *How an Algorithm Kicks Small Businesses Out of the Food Stamps Program on Dubious Fraud Charges*, COUNTER (Oct. 8, 2018), <https://thecounter.org/usdaalgorithm-food-stamp-snap-fraud-small-businesses/>); Chris McGann, *Somali Grocers Lose Right to Use Food Stamps*, SEATTLE PI (Apr. 8, 2002), <https://www.seattlepi.com/news/article/Somali-grocers-lose-right-to-use-food-stamps-1084746.php/>.

some explanation of why an agency believes its dataset covers a comprehensive range of test cases. Merely requiring an agency to explain why it does not believe the long-tail problem poses an issue will make the agency acknowledge the possibility of the problem. In addition, agencies might also seek the input of potentially affected parties, such as members of disadvantaged groups, to make sure the dataset covers any potential blind spots.<sup>191</sup> During notice and comment, such parties might raise concerns about how an algorithm will perform for subsets of the population. The reasons an agency gives in response to these concerns would likely entail some assessment of the performance of the algorithm for members of these groups, thus mitigating concern about the long-tail problem and improving the quality of the algorithm.

Beyond the long-tail problem, the data within the training dataset might also just be poor quality. Humans produce the data that goes into machine learning algorithms; if that data is poor quality—whether inaccurate, biased, or inconsistent—the algorithm is liable to be poor quality as well. As the adage goes, “garbage in, garbage out.”<sup>192</sup> With a policy, at least an agency can reason from first principles about what a quality decision should look like. If an agency decides to train a machine learning algorithm, however, it probably will rely at least in part on data from prior decisions. An agency should therefore scrutinize its data, lest it risk incorporating low-quality decisions by individual adjudications into an agency-wide policy.

Given the data-quality issue, agencies building algorithms should give reasons for believing their dataset meets standards for quality. The sheer volume of data used to train machine learning algorithms makes providing reasons difficult. And unlike other forms of reason-giving, for which public involvement is possible, agencies might not want to expose datasets to the public due to privacy concerns. Moreover, explaining why every datapoint in the dataset is of high quality would be impossible. Nevertheless, the agency could at least show that the algorithm performed as expected against a reasonably large number of test cases and explain how it scrutinized those test cases. David Freeman Engstrom and Daniel

---

191. See Nicol Turner Lee & Paul Resnick, *Algorithmic Bias Detection and Mitigation: Best Practices and Policies to Reduce Consumer Harms*, BROOKINGS (May 22, 2019), <https://www.brookings.edu/research/algorithmic-bias-detection-and-mitigation-best-practices-and-policies-to-reduce-consumer-harms/> (“Getting users engaged early and throughout the process will prompt improvements to the algorithms, which ultimately leads to improved user experiences.”).

192. One scholar has gone further, arguing that in an unequal society, the very act of prediction will reproduce that inequality. See Sandra Mayson, *Bias In, Bias Out*, 128 YALE L.J. 2218, 2224 (2019) (internal quotes omitted).



Ho call this process “prospective benchmarking.”<sup>193</sup> The reason-giving required by prospective benchmarking would, in at least one respect, make reason-giving for algorithms better than systemic reason-giving for policies, since agencies do not always have to consider a large number of test cases when designing policies.

Another way agencies could justify their use of a particular dataset without exposing the underlying data is to explain how the agency chose to include certain factors in the dataset. One way these explanations would promote quality decision making is by surfacing proxies for discrimination. Algorithms pose bias concerns even when they are not explicitly trained on protected factors such as race and gender. That is because protected factors shape people’s experiences of the world, making more seemingly neutral factors—zip code, insurance status, income—reliable proxies for protected categories. For some factors, like zip codes, the relationship to other protected categories might be obvious. An agency could certainly justify using such a factor, perhaps by pointing to accuracy improvements when the agency includes it—but such an explanation would at least lead the agency to consider potential bias concerns.

But other factors might be less obvious as proxies for protected categories. For example, only after extensive research did developers of a health care risk algorithm learn that healthcare costs tracked race.<sup>194</sup> Both algorithms and policies present the risk of hidden proxies. The difference is that humans probably choose the factors weighed in a policy based on their intuitions or knowledge about the world, while the factors on which an algorithm is trained might have just been one of many factors included in a pre-existing dataset, chosen for no reason other than that the agency already recorded the information. However, if an agency were required to give reasons for including certain factors, the agency might learn through research and due diligence that a factor has a possible relationship to protected categories. In justifying the choice of dataset, an agency may also want to explain where the data came from. Knowing the provenance of a dataset can be important because social inequalities can affect what gets measured and from whom.<sup>195</sup> In thinking about where

---

193. See Engstrom & Ho, *supra* note 56, at 849 (“The core idea is that when agencies adopt an AI decision making tool, they should subject it to benchmarking relative to a random hold-out set of cases that undergo conventional human review.”).

194. See Ziad Obermeyer, Brian Powers, Christine Vogeli & Sendhil Mullainathan, *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, 366 *SCI.* 447 (2019).

195. See, e.g., Joanna Radin, “Digital Natives”: *How Medical and Indigenous Histories Matter for Big Data*, 32 *OSIRIS* 43 (2017) (detailing the creation and circulation of the Pima Indian Diabetes Dataset, which was created based on the health information of Indigenous participants in a National Institutes of Health study on diabetes and was later used to refine algorithms having nothing to do with diabetes or health).

data came from, agency staff might recognize factors that owe their predictive power to bias, arbitrariness, or otherwise undesirable causes.

So far, this Article has discussed the systemic reasons an agency gives before deploying an algorithm. But remember that reason-giving is an ongoing process. Auditing of algorithms' performance can help agencies generate robust systemic reasons for subsequent versions of an algorithm. The National Health Service's audit of the AFST is a case in point. As explained in Section II.A, auditing—spurred by public criticism—led the DHS to change the target outcome of its algorithm as well as the sources of data going into it. The agency gave reasons for these changes grounded in data obtained from audits of the first version of the algorithm.<sup>196</sup>

One general concern that looms over algorithmic reason-giving is that the systemic reasons might be hard for non-technical audiences to understand. The following Part focuses on these technical gaps. For now, the important point is that algorithms could undermine quality relative to policies because some of the reasons might be intelligible only to a small set of stakeholders. Agencies have taken steps to address this problem, such as by increasing the number of technical employees in the government.<sup>197</sup> To the extent possible, however, agencies should also try to explain technical decisions with non-technical reasons.

Accordingly, systemic reason-giving for algorithms promotes quality by encouraging deliberation, engagement with different stakeholders, and, most obviously, *reasoning*. Moreover, with a more explicit focus on error rates, actual case outcomes, and performance over time, algorithmic reason-giving might do more to promote quality decision making than systemic reasons for policies. The biggest challenges relate to the quality of the underlying data and to gaps between technical and non-technical stakeholders. Reasons that make sense to one group might not make sense to another, and vice-versa. Whether systemic reasons for algorithms promote quality, then, will depend in part on how agencies structure the teams responsible for building the algorithms. A mix of domain expertise and technical acumen is ideal.

In the end, the important question might not be whether quality-promoting systemic reason-giving for algorithms is possible, but whether it is feasible. The DHS spent years studying the AFST, enlisting the help of researchers from top universities. And to build its algorithms, the SSA hired a team of technically adept lawyers and engineers. In theory,

---

196. See *supra* Section II.A.

197. See *FACT SHEET: Biden-Harris Administration Announces New AI Actions and Receives Additional Major Voluntary Commitment on AI*, <https://www.whitehouse.gov/briefing-room/statements-releases/2024/07/26/fact-sheet-biden-harris-administration-announces-new-ai-actions-and-receives-additional-major-voluntary-commitment-on-ai/> (July 26, 2024) (explaining that AI experts “are working on critical AI missions, such as informing efforts to use AI for permitting, advising on AI investments across the federal government, and writing policy for the use of AI in government.”).

agencies can take many steps to generating quality-promoting reasons. But will they have the technical expertise, time, and financial resources to do so?

## 2. Fostering Accountability

When agencies give reasons for algorithms, different audiences can scrutinize them and make sure the agency's reasoning is sound. However, the technical nature of algorithms and of the systemic reasons given for them potentially undermine accountability. This problem is not distinctive to algorithms, as evaluating many policies require expertise in the policy area. However, the problem is inevitable to a degree that it is not for policies. An important question, then, is who will be able to hold agencies accountable.

At least some systemic reasons for any algorithm will be technical. To start with an extreme example, consider an algorithm the EPA has considered using to identify chemicals that perform specific functions within the body.<sup>198</sup> The algorithm is part of the EPA's Exposure Forecasting (ExpoCast) project, which aims to generate exposure predictions so the agency can promulgate better rules under the Toxic Substances Act.<sup>199</sup> Those who built the algorithm justified their method for reducing the number of factors in the dataset to a manageable level as follows: "The method was chosen because it is a cursory approach to high-throughput categorization, which can be easily automated, and thus allow thousands of chemicals to be incorporated into the FUse dataset."<sup>200</sup> Those hoping to hold the EPA accountable would need some expertise in machine learning to respond to this reason, possibly through expert consultants.

To at least some degree, then, holding agencies accountable for their algorithms will involve a battle of the experts. Not all reasons for an algorithm need to facilitate the participation of any and everyone. Nor will all reasons for an algorithm be technical. As explained above, some

---

198. See Katherine A. Phillips, John F. Wambaugh, Christopher M. Grulke, Kathie L. Dionisio, & Kristin K. Isaacs, *High-Throughput Screening of Chemicals as Functional Substitutes Using Structure-Based Classification Models*, 19 GREEN CHEMISTRY 1063 (2017).

199. See *Rapid Chemical Exposure and Dose Research*, Env't Prot. Agency (2018), [https://www.epa.gov/sites/default/files/2014-12/documents/exposure\\_forecasting\\_factsheet.pdf](https://www.epa.gov/sites/default/files/2014-12/documents/exposure_forecasting_factsheet.pdf). The agency primarily uses the algorithm to support research. But one could imagine the agency promulgating a rule that subjects all chemicals that achieve a certain score to increased scrutiny. This would be like the EPA's practice of regulating polluters who introduce more than a prescribed volume of lead into paint. Instead of a certain volume of pollutant, an algorithmic score would be the determining factor in the policy. See *Hazard Standards and Clearance Levels for Lead in Paint, Dust and Soil (TSCA Sections 402 and 403)*, Env't Prot. Agency, <https://www.epa.gov/lead/hazard-standards-and-clearance-levels-lead-paint-dust-and-soil-tsc-sections-402-and-403>.

200. See Phillips et al., *supra* note 198, at 1072.

reasons for algorithms are grounded in policy priorities, and many technical choices can be justified in non-expert terms. But this is cold comfort, as the most technically complex reasons may be the ones justifying crucial aspects of the algorithm's design. From an accountability perspective, then, systemic reasons for algorithms seem to be in trouble.

But remember that the task is comparative. Technically and scientifically complex reasons are not new to algorithms. One of the primary reasons Congress delegates policymaking power to agencies is because they possess specialized expertise.<sup>201</sup> Policies, and the reasons given for them, have become more technically and scientifically complex over time.<sup>202</sup> A clear dividing line between the science and politics of policy remains elusive, and in practice, the two are intricately intertwined.<sup>203</sup> Whether designing algorithms or policies, agencies cannot articulate every reason into terms that non-experts will understand. Sometimes, that might mean that agencies can deliberately misuse its scientific or technical expertise to hide unpopular choices and receive more deferential judicial review.<sup>204</sup> More specifically, agencies may emphasize the scientific reasons for policy choices that are, at bottom, value judgments.<sup>205</sup> For example, when revising the ozone standard under the Clean Air Act, EPA scientists could not reach a consensus on a single quantitative standard. The final standard was a compromise between White House concerns for the economy on the one hand and public health concerns on the other. However, in its published explanation in the Federal Register, the EPA gave exhaustive scientific reasons for the standard, numbering some fifteen pages.<sup>206</sup> The gambit worked, and the D.C. Circuit held that the EPA had arrived at the standard taking "into account all the relevant studies . . . in a rational manner . . ."<sup>207</sup>

So, while the technical nature of systemic reasons for algorithms does pose problems for accountability, these problems are not new. A full

---

201. See Susan Webb Yackee, *The Politics of Rulemaking in the United States*, 22 ANN. REV. OF POL. SCI. 37, 40 (2019).

202. See Wendy E. Wagner, *A Place for Agency Expertise: Reconciling Agency Expertise With Presidential Power*, 115 COLUM. L. REV. 2019, 2023 (2015) ("[T]he hypertechnicality of agency rules is a more recent phenomenon . . .").

203. WENDY WAGNER, SCIENCE IN REGULATION: A STUDY OF AGENCY DECISIONMAKING APPROACHES 16 (2013), [https://www.acus.gov/sites/default/files/documents/Science%20in%20Regulation\\_Final%20Report\\_2\\_18\\_13\\_0.pdf](https://www.acus.gov/sites/default/files/documents/Science%20in%20Regulation_Final%20Report_2_18_13_0.pdf).

204. See Emily Hammond Meazell, *Super Deference, the Science Obsession, and Judicial Review as Translation of Agency Science*, 109 MICH. L. REV. 733, 735-36 (2011) ("[Agency science] is laced with policy decisions at numerous levels, making it susceptible to misuse.").

205. See Wendy E. Wagner, *The Science Charade in Toxic Risk Regulation*, 95 COLUM. L. REV. 1613 (1995).

206. *Id.* at 1640-44 (explaining the process of revising the ozone standard).

207. See *American Petroleum Inst. v. Costle*, 665 F.2d 1176, 1187 (D.C. Cir. 1981).

defense of agencies' reliance on scientific expertise is outside the scope of this Article. Suffice it to say that agencies could not do their jobs without offering scientific or technical reasons for their policies. Still, some degree of accountability is always possible. For many rules, the notice and comment stage is already largely the domain of well-resourced parties. Just as these parties can retain scientific experts, they can also retain machine learning experts. The same goes for parties challenging agency rules in court. In addition, judges can serve as translators, using their opinions to convert an agency's stated reasons for an algorithm into terms the public can grasp, just as they have done for policies.<sup>208</sup>

What is distinctive about algorithmic reasons is that they inject technical reasons into domains where they may not have been present before. It is quite a different matter to introduce an algorithm to the EPA—which already relies on complex scientific models—than it is to introduce one to a less technical agency. A distinction between technical and domain expertise is useful here. Understanding the workings of an agency like the Department of Homeland Security no doubt requires a tremendous degree of expertise. But technical expertise—namely, expertise in machine learning or similar methods—is likely less critical. Machine learning methods dress domain expertise in technical attire, making technical expertise important to a degree that it likely was not before. That is one way in which systemic reason-giving for algorithms might promote accountability less than systemic reason-giving for policies.

### 3. Preserving Dignitary Interests

Systemic reasons for algorithms protect dignity in largely the same ways as systemic reasons for policies. They turn lawmaking into a conversation, signal concern for the regulated subject's interests, and facilitate contestation. These effects evince an agency's respect for interested parties. Yet, the scientific and technical complexity of reasons decreases their dignity-protecting value. Contestation becomes more difficult as technical complexity rises. Moreover, technical complexity may make it more difficult for regulated parties to know whether the agency is meaningfully responding to their concerns. However, systemic reasons for algorithms are mostly indistinguishable on dignitary grounds from other technically or scientifically complex reasons. Just as it would be impossible for agencies to completely avoid scientifically complex reasons out of concern for dignity, it is also untenable to reject algorithms entirely because some reasons given for them are technical. But where

---

208. See Hammond Meazell, *supra* note 204, at 778 (describing the role of courts as translators).

giving entirely non-technical reasons is possible—in other words, where it is possible to have a non-technical policy instead of an algorithm—then choosing to use an algorithm leads to dignitary costs. These costs need not be dispositive, but they do tip the balance in favor of using a policy instead of an algorithm, all else equal.

### III. CASE-SPECIFIC REASONS (OR ALGORITHMS/POLICIES AS REASONS)

Usually, when we say that algorithms are opaque, we mean that algorithms do not explain why a particular combination of inputs led to a particular output.<sup>209</sup> That is why commentary largely focuses on the black-box problem. While systemic reasons shed light on how algorithms operate, fixing one's attention on false positive and negative rates is different from attending to the circumstances of an individual case. An individual denied benefits in part because of an algorithm might ask why that was so. Pointing to the algorithm's overall accuracy would feel like a misdirection. The person denied benefits would probably still feel like they never got an answer to the question, "why me?"

This Section is about the "why me" question. Denying that algorithms undermine reason-giving at the individual level is a losing game. Some commentators have taken a different tack, highlighting the promise of so-called explainable AI, which includes algorithms that purport to offer insight into individual decisions.<sup>210</sup> But such algorithms comprise, at most, a small minority of those being used by agencies, and the explanations given by these algorithms often fall short of fully individualized reasons for a decision—meaning explainable AI may be a false hope.<sup>211</sup>

Instead of pinning its hopes on new technological advances, this Section, like the one before it, uses the algorithms-as-policies framework.

---

209. See Burrell, *supra* note 22, at 1.

210. See Boris Babic & I. Glenn Cohen, *The Algorithmic Explainability "Bait and Switch"*, 108 Minn. L. Rev. 857, 862-63 (gathering sources that have claimed that explainable AI is more trustworthy, easier to understand, safer, and more accountable/transparent). Babic and Cohen conclude that "explainable AI/ML models fundamentally fail to achieve these goals: these models fail to assist users in both correctly interpreting a model and in understanding the true reasons or principal factors behind the model's predictions." *Id.* at 864.

211. See *id.*; Marzyeh Ghassemi, Luke Oakden-Rayner & Andrew L. Beam, *The False Hope of Current Approaches to Explainable Artificial Intelligence in Health Care*, 3 THE LANCET DIGIT. HEALTH e745, e745 (2021) ("In practice, explanations can be extremely useful when applied to global AI processes, such as model development, knowledge discovery, and audit, but they are rarely informative with respect to individual decisions."); Jeremy Kahn, *Software Vendors are Pushing "Explainable A.I." that Often Isn't*, FORTUNE (Mar. 22, 2022), <https://fortune.com/2022/03/22/ai-explainable-radiology-medicine-crisis-eye-on-ai/> ("Everyone who is serious in the field knows that most of today's explainable A.I. is nonsense . . . ." (quoting Zachary Lipton, Professor of Computer Science, Carnegie Mellon University)).

It inquires into what extent algorithms undermine reason-giving at the level of a single case, and how that compares to the effect that policies have on case-specific reason-giving. In answering that question, this Section shows that algorithms reveal a longstanding tension in administrative governance between bureaucratic management and individualized consideration. Recognizing that algorithms did not invent this tension, this Section provides a way to resolve it. In short, the argument is that algorithms, like policies, can be valid reasons for agency action *in and of themselves*.

#### *A. An Old Problem Revisited: Case-Specific Reason-Giving*

Like algorithms, policies undermine individualized reason-giving. Consider how a frontline adjudicator would use a policy like the medical-vocational grid.<sup>212</sup> First, the adjudicator would make factual determinations about the claimant's age, education, and previous work experience. Then, the adjudicator would look to the grid, which would direct a finding of disabled or not disabled based on these determinations. Once the grid is in place, it is no longer the adjudicator's job to give reasons for making a disability determination based on the factors considered. The adjudicator also does not need to explain why they relied on age, education, and previous work experience to make their determination, as the adjudicator is statutorily required to consider these factors.

In reducing the number of issues over which the frontline adjudicator has discretion, the grid narrows the number of decisions for which the adjudicator must provide reasons. Granted, the grid regulations allow adjudicators to make exceptions.<sup>213</sup> And under some circumstances, the grid does not apply.<sup>214</sup> But while the adjudicator may give reasons for granting an exception, as the grid allows, they need not re-justify using the policy every time they apply it. Doing so would defeat the efficiency and managerial grounds for implementing a policy in the first place. Requiring a decision maker to re-justify a binding policy, such as a legislative rule without exceptions, would be more than inefficient: it would be a contradiction in terms. Whether the decision maker had reasons for the policy or not, they would have to apply it.

---

212. For discussion of the medical-vocational grid, see *supra* Section I.D.1.

213. See, e.g., Dubin, *supra* note 120, at 942 ("The broadest and most litigated exception to the grid's direct application is in situations involving claimants with nonexertional or non-strength related medical limitations.")

214. See, e.g., Hall v. Chater, 62 F.3d 220, 224 (8th Cir. 1995) ("If, however, the claimant also has a non-exertional impairment, such as pain, the Secretary must use vocational expert testimony or other similar evidence to meet the burden of showing the existence of jobs in the national economy that the claimant is capable of performing.") (internal quotation marks omitted).

With a policy or algorithm in place, the decision maker does not have to re-invent the wheel. The policy or algorithm is a shortcut. But in taking that shortcut, the decision maker's reasons for reaching a result will be less tied to the particular case at hand than they would have been absent the policy or algorithm. Consider a hypothetical world in which adjudicators made disability determinations without resorting to any policy. In that world, the adjudicator could consider any factor in making a determination. Therefore, the adjudicator may need to give reasons for relying on some factors but not others. After settling on the important factors, the adjudicator would need explain why a specific combination of factors supports a determination of disabled or not disabled. Without policies in place, such explanations could appeal to a limitless number of general principles—perhaps that older people who cannot stand comfortably for long periods at a time are properly deemed disabled, or that teenagers who live in a particular region are unlikely to find work if they suffer from a visual impairment. Where decision makers make decisions without resort to any policy or algorithm, such principles could be tied to the specific circumstances of an individual claimant. For example, if a claimant from a rural village that offers residents only manual labor jobs cannot stand for long periods of time, the adjudicator may deem that person disabled based on the narrow principle that residents of that particular village are disabled when they cannot stand for long periods of time.

Giving a reason is always an act of generalization. A reason connects particular facts to an outcome by appealing to some principle more encompassing than the case at hand.<sup>215</sup> The more specific the principle, the less the reason resembles a reason and the more it resembles a mere description of the case. Only where the principle undergirding the reason is sufficiently general as to apply to a large range of cases do decisions become more predictable. But the more general the principle, the less the decision is grounded in the circumstances of the individual claimant. That is because committing to making decisions based on a general principle means agreeing in advance to exclude potentially relevant facts about a claimant from consideration.<sup>216</sup>

It was precisely this potential exclusion of relevant facts that animated much of the public's concern about the medical-vocational grid.<sup>217</sup> One commentator stated explicitly that “the proposed regulations provide an excuse for the SSA not to gather all facts pertinent to an individual.”<sup>218</sup> Similarly, others worried that the grid would lead the agency to “make

---

215. See Frederick Schauer, *Giving Reasons*, 47 STAN. L. REV. 633, 635 (1995).

216. *Id.* at 651.

217. See MASHAW, *supra* note 52, at 117-20.

218. *Id.*



determinations based on the ‘average man’ concept” and would lead the agency to “treat[] claimants according to categories or classes rather than as individuals . . . .”<sup>219</sup> Rephrasing these comments with a view towards reason-giving, a large concern was that the grid would lead decision makers to make decisions based on reasons that were inattentive to a person’s circumstances.

The point is that policies, like algorithms, circumscribe individualized reason-giving. Both limit the possible answers to the question “why me”—sometimes so much that the answer amounts to little more than that “the policy or algorithm said so.” Systemic reasons do not entirely solve this problem because systemic reasons are just that—reasons for an overall system design, not for deciding a particular case in a certain way. In this way, the problem policies and algorithms pose for individualized reason-giving reflects a core tension in administrative law between bureaucratic management and frontline discretion. Indeed, the extent to which the policy or algorithm limits individualized reason-giving depends on the degree to which the policy or algorithm binds the agency employees applying it.

Where decision makers have wider latitude to deviate from a policy or algorithm—and thus where bureaucratic management has a lighter touch— decision makers have greater flexibility to consider the unique circumstances of a particular case. In turn, decision makers have greater flexibility to give individualized reasons for a decision that are based on these unique circumstances.<sup>220</sup> In contrast, where policies or algorithms are binding—where bureaucratic management is at its most restrictive— reasons the decision maker independently gives for the decision are mere post hoc rationalizations. The answer to “why me” in that case is that bureaucratic managers decided in advance, without reference to the particular facts about the individual party, that it should be so. The question then becomes, should we tolerate that?

### *B. A Way Out: Policies/Algorithms as Reasons*

To at least some extent, the reason a person gives when making a decision based on an algorithm or policy will take the form: “because the algorithm/policy said so.” The policy, on its face, may give indications of the underlying intuitions behind the decision, but the reasons behind a policy are not up for debate when the policy is applied to a particular case.

---

219. See Rules for Adjudicating Disability Claims, *supra* note 113, at 9291.

220. Relatedly, where a policy resembles a standard rather than a rule, the frontline decision-maker has greater latitude to consider the unique circumstances of a particular case and to give reasons accordingly.

Consider a speed limit.<sup>221</sup> A city may set a sixty-five mile per hour speed limit on a particular stretch of road because doing so promotes safe driving while recognizing people's need to reach destinations quickly. However, these reasons may not apply in some traffic conditions. It may be more dangerous, for example, to drive sixty-five miles per hour than seventy-five miles per hour if the flow of traffic is moving at eighty miles per hour. Nevertheless, a traffic enforcement officer can give a ticket to someone driving seventy-five miles per hour, even in the eighty mile per hour flow of traffic. In doing so, the officer does not have to justify applying the policy anew by giving reasons for thinking that driving seventy-five miles per hour is unsafe. That work was done on the front end by regulators or legislatures that set the speed limit. When the officer is writing a ticket, the speed limit *is* the reason.

That is how to understand policies and algorithms at the case level: as reasons in and of themselves. A reason connects facts to an outcome by reference to a general principle. The policy or algorithm supplies that principle. As Joseph Raz has most clearly explained, policies are reasons in two ways.<sup>222</sup> First, they are reasons to carry out some action. For instance, when the SSA deems a claimant disabled, it can give the medical-vocational grid as reason for doing so. Second, they are reasons *not* to act for competing reasons. In this sense, they are exclusionary. Using the same example, the medical-vocational grid gives the agency a reason not to consider factors beyond a claimant's age, medical status, and work experience. At this most basic level, algorithms act as reasons, too. When the DHS investigator investigates a particular family because the AFST flagged it, the algorithm is acting as a reason to carry out an investigation. Moreover, the AFST is also a reason to exclude other reasons to act differently.

To say policies or algorithms are reasons is not to say they are necessarily outcome determinative. In his discussion of agency guidance—a kind of policy—Blake Emerson borrows from Raz to argue that guidance has legal authority insofar as it offers privileged reasons for agency action.<sup>223</sup> The guidance is a privileged reason because, while not necessarily dispositive, the positions it states cannot be dismissed lightly.<sup>224</sup> Guidance is thus a rule of thumb that gives the agency some standard to apply, without precluding the possibility that the agency might depart from the guidance if countervailing reasons make doing so

---

221. This example comes from FREDERICK SCHAUER, *PLAYING BY THE RULES: A PHILOSOPHICAL EXAMINATION OF RULE-BASED DECISION-MAKING IN LAW AND IN LIFE* 95-96 (1993).

222. See Raz, *supra* note 54, at 216.

223. See Blake Emerson, *The Claims of Official Reason: Administrative Guidance on Social Inclusion*, 128 *YALE L.J.* 2122, 2133 (2019).

224. *Id.* at 2149.

desirable.<sup>225</sup> The important point is that where a policy or algorithm is non-binding, the policy or algorithm is a privileged reason for making a certain decision. Only where the policy or algorithm is binding does it cancel out other reasons for acting differently. In addition, unless a policy or algorithm is binding, it need not be the sole reason for making a certain decision. The decision maker could give a host of other reasons for taking certain action, but the policy or algorithm would remain a presumptively valid reason for acting.

Compared to other kinds of reasons, policies and algorithms are distinctive. Typically, a reason shows what is good or right in an action—I took out the trash because it makes my house cleaner; I invited Sarah because she is a kind person. Policies and algorithms do not necessarily reveal what is good, only what must be done.<sup>226</sup> In some cases, we can infer what is good and what must be done based on the policy or algorithm. For example, when the EPA imposes an emissions standard, we can infer that action taken according to it probably promotes environmental goals in some way, and likely without excessively burdening industry. These inferences are based not on the policy on its face but in a broader understanding of how the EPA operates and what it typically hopes to achieve. Similarly, when the DHS carries out a screening based on the AFST, our knowledge of the agency suggests that the screening is meant to promote child safety. That said, the evaluative character of the policy or algorithm is not at the core of what makes it a valid reason.<sup>227</sup>

How can policies and algorithms be justified as reasons for agency action? In the typical statement of the black-box problem, the issue is that when one asks why the algorithm produced a particular result based on the inputs, the most direct answer, without reference to systemic reasons, is: “because the algorithm said so.” This problem is less troubling than it may initially seem if the algorithm is a valid reason for agency action. At this point, it is worth remembering that this Article aims only to evaluate algorithmic reason-giving relative to policy reason-giving. So long as algorithms do not fall too short, as reasons, compared to policies, they remain defensible.

An instrumentalist justification for policies as reasons immediately presents itself. Under such a view, the purpose of a policy is precisely for it to serve as a reason for action. If policies lacked force as reasons, agencies would have no reason to issue them. But such instrumentalism does little to justify algorithms because algorithms are a relatively novel

---

225. *Id.* at 2150.

226. See Binesh Hass, *The Opacity of Rules*, 41 OXFORD J. OF LEGAL STUD. 407, 407-08 (2021).

227. *Id.*

tool of government. While agencies create algorithms hoping for them to act as reasons, there is little historical basis to believe members of the public or the courts will treat them as such. The same goes for algorithms used internally by agency employees. Given skepticism towards algorithms as reasons, one should move beyond instrumentalism and consider more precisely what makes policies and algorithms valid reasons for agency action.

This Article's view is that policies and algorithms are valid reasons in large part because of the systemic reasons given for them. The law reflects this view: Agencies are generally required to give systemic reasons for policies that are binding on regulated parties, and systemic reasons comprise the record on which courts judge the legality of agency rules.<sup>228</sup> Notice and comment is the procedure that makes agency rules legally valid as reasons for agency action; systemic reasons are the language of the notice and comment process.

Systemic reasons matter when judging whether an algorithm or policy should serve as a reason, but it is important to be clear about why they matter. Grounding the validity of policies-qua-reasons in the systemic reasons given for them makes intuitive normative sense. Reasons are important because they promote decisional quality, foster accountability, and preserve dignitary interests.<sup>229</sup> The validity of policies or algorithms as reasons should depend, at least to some extent, on whether they further these objectives. Systemic reasons can help them do so.<sup>230</sup>

Still, a disconnect remains. While systemic reasons help policies or algorithms meet these objectives, they do so at the systemic level. Zooming in on the individual case, policies and algorithms sometimes do precisely the opposite. For one, the designers of the policy or algorithm might have had in mind different fact patterns from the one in the case at issue, making the policy or algorithm inapt for the case at hand. For another, if a policy or algorithm is binding and has properly gone through notice and comment, the regulated party in the individual case can no longer hold the decision maker accountable for the outcome. Most worrying, dignitary concerns remain pressing in the individual case, notwithstanding the policy or algorithm. "Because the policy or algorithm said so" is not a reason that is responsive to the concerns of the individual, and listing reasons for designing the policy or algorithm in a particular way can only go so far towards demonstrating solicitude for the individual regulated party. At root, the issue is the degree to which agencies focus

---

228. See *Auto. Parts & Accessories Ass'n v. Boyd*, 407 F.2d 330, 336 (1968) ("[T]here is a record compiled in a Section 4 proceeding, and available for filing in court. It consists of the submissions made in response to the invitations issued for written comments.")

229. See *supra* Section II.B.

230. See *supra* Section II.C.

on systemic justice over individual rights (though these are not necessarily at odds). Reason-giving is epiphenomenal.

Thus, one virtue of the algorithms-as-policy view is that it reveals when we are wading into old, intractable debates. The promulgation of the medical-vocational grid was a flashpoint because it cut to the heart of an unavoidable tension in administration between individual consideration and systemic rationality.<sup>231</sup> Reasonable minds will disagree about when decision making should be left to individual bureaucrats and when agencies should control decision making from above. One may believe individualized consideration is more important than systemic accuracy or efficiency, or may be skeptical that, in certain domains, decisions can be made more accurate or consistent through policies or algorithms. Wherever one may draw a line in the sand between acceptable and unacceptable limits on individualized consideration—and one must draw this line somewhere—the claim here is only that the line for policies and the line for algorithms should be drawn relatively close together.

Returning to systemic reasons, these are important to the validity of algorithms and policies as reasons, but not when viewed from the level of the individual case. To see why they matter, it is necessary to zoom out and ask why agencies countenance using policies or algorithms in the first place, particularly given the effect they have on individualized reason-giving.

Essentially, policies are valid as reasons for agency action because they are necessary for administration. Administration requires a balance of accuracy, efficiency, predictability, and accountability. Agencies therefore need tools like policies to achieve these goals. Were policies not to serve as authoritative, or at least persuasive, reasons for action, agencies would struggle to carry out their priorities, for everything would depend on the discretion of the individual decision maker. So, even accepting the possible deficiencies of policies as reasons—namely, that they do not necessarily reveal what is good in agency action—policies are defensible as reasons because they promote desirable features of a bureaucracy. Similarly, algorithms can serve as reasons so long as they also help the agency carry out these same bureaucratic goals.<sup>232</sup>

Certainly, policies or algorithms are not always valid as reasons just because they serve certain bureaucratic goals. One can imagine a range of scenarios that are too fact-specific, or just too sensitive, to be adequately served by a policy or algorithm, at least one that takes considerable discretion away from an individual decision-maker. In some circumstances, the costs to individualized reason-giving may outweigh

---

231. *See supra* Section I.D.1.

232. *See supra* Section II.C.

systemic benefits of other kinds. Systemic reasons are important to the validity of algorithms and policies as reasons because they show how the agency considered these important tradeoffs. Through systemic reasons, agencies can explain why they believe a policy or algorithm will promote important bureaucratic objectives. Whether agencies realize it or not, they are thereby explaining why the policy or algorithm should be a valid reason for agency action.

Perhaps surprisingly, then, the validity of a policy or algorithm as a reason does not primarily depend on the policy or algorithm itself. Instead, their validity as reasons depends on the justifications given for them. Admittedly, these things are closely related. But the distinction matters because it makes less pressing a problem that seems to plague algorithms acting as reasons, which is that they are not intuitive on their face. One cannot read an algorithm in the way one can read, say, the medical-vocational guidelines. Looking at the algorithm alone, it is impossible to develop intuition about whether the algorithm is a sensible reason for agency action. For algorithms used externally, this problem raises dignitary concerns (the regulated party can't intuitively grasp why the agency made a particular decision about them), quality concerns (one cannot know whether a decision *feels* correct, perhaps to surface problems with the algorithm that merit attention), and accountability concerns (it is hard to contest an algorithm in the individual case if one cannot intuit its logic). One should not overestimate the degree to which policies are intuitive on their face, whether to employees or regulated parties. Yet, even granting that algorithms are less intuitive on their face than policies, that is not necessarily a reason to strongly prefer policies to algorithms. Instead, what matters more are the systemic reasons given for the algorithm or policy.

#### IV. IMPLICATIONS

Algorithmic and policy reason-giving encompass two different processes: the giving of systemic reasons for the algorithm or policy, and the giving of the algorithm or policy as a reason when deciding a particular case. Algorithms indeed pose more problems for both processes than policies do. Many of the systemic reasons for algorithms are technically complex, and algorithms are less intuitive on their face as reasons than policies. These deficiencies are best thought of as weights on the scale that might lead an agency to opt for a policy rather than an algorithm.

The task ahead is twofold: first, it is to suggest how an agency can decide whether an algorithm is adequate on the reason-giving front. To that end, this Section offers a normative framework that agencies and

courts can use to assess whether to use an algorithm or a policy for a particular task. A second task is to evaluate whether the law reflects this Article's normative conclusions. In that vein, this Section concludes that agencies should be able to use at least some algorithms without running afoul of the arbitrary and capricious standard.

#### *A. Normative Framework*

Imagine a policy and algorithm designed to carry out the same task. Assume the policy and algorithm are relatively equal in their expected accuracy, consistency, and efficiency. Ignoring other potential considerations, would limitations in reason-giving be grounds for choosing the policy over the algorithm? In some instances, yes. When it comes to systemic reasons, the inherently technical nature of algorithms posed challenges for quality, accountability, and dignity. As for case-specific reasons, the intuition gap (i.e., machine learning algorithms are not intuitive on their face) likewise raises concerns about dignity, quality, and contestation. All else equal, these two problems typically will tip the scales in favor of policies.

Yet, understanding the systemic reasons for a policy, and understanding the policy as a reason, can also be difficult. Any policy—indeed, any reason—sits somewhere on one continuum of technicality and on another continuum of intuitiveness. It makes sense, then, that in areas such as environment and health policy, influential participants are typically sophisticated interest groups who can hire technical experts to help them hold agencies accountable.<sup>233</sup> In areas where technical sophistication is already part and parcel of agency action, the case for favoring policies over algorithms is weaker because the reason-giving gap between them is narrower. Similarly, where existing agency policies are already unintuitive on their face, replacing them with algorithms won't have a great effect on dignity, quality, or accountability. There, factors other than reason-giving should be dispositive.

Administration is all about trade-offs. One need not subscribe to crude utilitarianism to recognize as much. Even for a person committed to certain inviolable procedural rights, it is far from obvious that algorithms' deficiencies in reason-giving should lead one to declare algorithms verboten while conceding that policies are an essential feature of bureaucratic management. If one goes, the other should, too. And if one stays, reason-giving does not furnish adequate grounds for keeping the other out.

---

233. See Thomas J. Hwang, Jerry Avorn & Aaron S. Kesselheim, *Life Cycle of Medical Product Rules Issued by the US Food and Drug Administration*, 39 J. HEALTH POL. POL'Y L. 751, 754-55 (2014).

Assuming there is some role for algorithms in administration, how should agencies and judges go about weighing their deficiencies in reason-giving against other considerations? Put differently, at what point does the inferior nature of algorithmic reason-giving mean the algorithm is no longer normatively justifiable? This inquiry is made more difficult by the fact that agencies rarely will have a fully developed policy alternative to an algorithm, meaning agencies cannot simply assess which is likely to be the most accurate, consistent, and efficient and then weigh reason-giving concerns accordingly.

What is needed is a framework for evaluating an algorithm based on the information the agency or court would have on hand—namely, the systemic reasons for the algorithm and the context in which it is used.<sup>234</sup> This Article proposes that agencies and judges should weigh four broad considerations in evaluating algorithmic reason-giving: (1) the systemic reasons given for the algorithm; (2) the intended recipient of the systemic reasons and the final algorithm's output; (3) the nature of the decision being made with the algorithm; and (4) the bureaucratic considerations that led the agency to try using an algorithm in the first place.

First, the agency, or court, should consider the systemic reasons given for the algorithm. As explained above, the systemic reasons given for an algorithm affect whether the policy is a valid reason for agency action. Broadly speaking, the better the systemic reason-giving, the more justifiable the algorithm. Algorithms that have undergone audits are typically preferable to those that have not, as systemic reasons surfaced through audits can shore up confidence that the algorithm will perform well in practice. As much as possible, agencies should aim to make the systemic reasons for an algorithm public. At a minimum, a broad range of stakeholders within the agency, both technical and non-technical, should have access to the systemic reasons to facilitate adequate contestation.

Second, the agency, or court, should consider the intended recipients of both the systemic reasons for the algorithm and the algorithm as a reason. If the algorithm is used purely internally, the agency should consider whether those privy to the systemic reasons are able to question them to an adequate degree. The same holds true if the systemic reasons are exposed publicly, perhaps through notice and comment. An important factor here will be the technical aptitude of such parties—or in the case

---

234. Other scholars have proposed frameworks to help agencies decide whether to use an algorithm. For instance, Cary Coglianese and Alicia Lai have developed a “meta-process” that can help an agency determine whether and when to use an algorithm as opposed to a human decision maker for a particular task. See Coglianese and Lai, *supra* note 45, at 1318. The framework developed in this Article is based mainly on the reasons given for an algorithm because of the centrality of reason-giving to administrative law.



of regulated parties, these parties' ability to retain technical experts. Whether the regulated party is a person or an entity is also important. The technical nature of systemic reasons for algorithms, as well as the unintuitive nature of algorithms as reasons, can have dignitary costs for regulated parties, as these parties will not necessarily understand why a particular decision was made about them. Entities such as corporations do not have pressing dignitary interests in the way that a disability claimant does.

Third, the agency, or court, should consider the nature of the decision being made using the algorithm. For one, how outcome-determinative is the decision? Is the algorithm, for example, being used as a core part of a final adjudication? In that case, the reason-giving costs of algorithms are particularly concerning. Unless the agency has good reason to believe, based on the systemic reasons available, that the algorithm will make decisions more consistent and accurate, the agency should do without the algorithm. The same is true if the decision being made is a sensitive one, with potentially life-altering consequences for the regulated party. In general, the less outcome-determinative and sensitive the decision, the less the agency should demand by way of systemic reasons. The reasons for designing an internal search algorithm in a certain way, for example, matter much less—and thus merit much less scrutiny—than the reasons for designing a disability adjudication algorithm.

Fourth and finally, the agency, or court, should consider the conditions that led the agency to decide to use an algorithm in the first place. If the current system is plagued by inconsistency, inaccuracy, and inefficiency, the costs to reason-giving should matter less than improving on the status quo. That is true whether the agency is considering creating an algorithm or a policy. Reasonable minds can disagree about how to weigh bureaucratic considerations against reason-giving, but at a minimum the quality of the former should have some influence on what will be tolerable in the latter.

This normative framework is important for both courts and agencies. Many algorithms, like many policies, will be shielded from judicial review.<sup>235</sup> While the use of algorithms, like the use of policies, may create avenues for judicial review where they may have otherwise not existed,<sup>236</sup> some underenforcement of algorithms will remain. Moreover, judicial review has significant disadvantages for addressing problems with agency action and is not necessarily an ideal way to root out systemic

---

235. See Engstrom & Ho, *supra* note 56, at 836 (“[C]onventional ex post judicial review of agency action under the APA is unlikely to generate systematic or even consistent review of the government’s new algorithmic toolkit in either the enforcement or adjudication context.”).

236. See Metzger & Stack, *supra* note 110, at 1281-83 (explaining how structured internal administration can authorize judicial review).

problems.<sup>237</sup> Accordingly, the onus will sometimes be on agencies to determine whether the algorithms they are using are adequate from a reason-giving perspective.

### *B. Arbitrary and Capricious Review*

In general, algorithms should survive arbitrary and capricious review based on the systemic reasons given for them, to the extent that policies survive the same. Section 706(2)(A) of the APA provides that courts shall set aside agency action that is “arbitrary [and] capricious.”<sup>238</sup> The reasons the agency gave for acting comprise the record for judicial review. For rulemaking, this means the reasons the agency gave during notice and comment, while in adjudications it means all the reasons the agency gave when making its decision. Depending on the context, then, arbitrary and capricious review of algorithms will turn on one of two inquiries. If litigants make a facial challenge to an algorithm promulgated as part of a rule, the issue will turn on the adequacy of the systemic reasons the agency gave for the algorithm. But if litigants make a challenge to a particular adjudication in which an algorithm is involved, the issue will be the extent to which the algorithm is an adequate reason for the outcome.

Arbitrary and capricious review is an open-ended inquiry, but patterns have emerged. To survive arbitrary and capricious review, agencies typically must respond to all comments,<sup>239</sup> consider the relevant factors as set forth by the substantive statute,<sup>240</sup> and consider reasonable alternatives.<sup>241</sup> Even under supposedly “hard look” review, the government wins most of the time, especially at the Supreme Court.<sup>242</sup>

If algorithms are fully ventilated during rulemaking, they should be able to survive arbitrary and capricious review. That is because, whether the rule takes the form of a policy or algorithm, the agency can still engage with the public and explain its thinking through systemic reasons. Just as an agency can respond to comments about a proposed policy, so it can respond to comments about a proposed algorithm. Moreover, agencies should be able to pass the “relevant factors” test when using algorithms by pointing to the factors that serve as the algorithm’s inputs. If

---

237. See Nicholas Bagley, *The Puzzling Presumption of Reviewability*, 127 HARV. L. REV. 1285, 1329-36 (2014) (discussing drawbacks of judicial review of agency action, including diverting agency resources, introducing delay, and limiting agency flexibility).

238. 5 U.S.C. § 706(2)(A) (2018).

239. See *Auto. Parts & Accessories Ass’n v. Boyd*, 407 F.2d 330, 341 (1968).

240. See *Motor Vehicle Mfrs. Ass’n v. State Farm Mut. Auto. Ins. Co.*, 463 U.S. 29, 42-43 (1983).

241. *Id.* at 51.

242. See Jacob Gersen & Adrian Vermeule, *Thin Rationality Review*, 114 MICH. L. REV. 1355, 1358 (2016).

commenters are concerned that the algorithm is not adequately accounting for certain inputs, agencies can vary the inputs to gauge the algorithm's sensitivity to the changes. It is still not a dead-end for the agency if the algorithm is insensitive to some factor, as the agency can then conduct research to explain why it thinks that factor is not outcome-determinative. Finally, nothing about algorithms precludes agencies from considering reasonable alternatives. Designing an algorithm is itself an exercise in weighing different alternative models against each other in an iterative manner.

Admittedly, this simple picture omits some difficulties. First, data disclosure is sometimes impossible due to ethical or legal considerations.<sup>243</sup> In this case, agencies should be required to give reasons to believe the data is accurate, perhaps by giving system-level statistics about it, and courts should weigh this deficiency using the normative framework above. Second, machine learning algorithms can change dynamically over time, in which case the agency should explain why it does not expect changes to be dramatic.<sup>244</sup> As long as the agency gives adequate responses to these concerns and shows that it has considered them, courts should refrain from second-guessing agencies' judgment. Possible difficulties notwithstanding, algorithms that go through notice and comment should mostly glide through arbitrary and capricious review under current law owing to the law's focus on the systemic reasons given for a rule.

However, algorithms that go through full notice and comment are the easy case.<sup>245</sup> They are also, in some sense, the less important case, as no algorithm in use by a federal agency has gone through notice and comment. What role can an algorithm play in agency decision making if it has not gone through notice and comment? In the context of arbitrary and capricious review, the issue turns on the degree to which such an algorithm can act as a reason for agency action.

For the agency to avoid putting the algorithm through notice and comment, the algorithm would need to be considered guidance in the form of a policy statement or interpretive rule. While some scholars have argued that agencies should have to defend guidance each time they use it if it does not go through notice and comment, effectively meaning that the guidance cannot act as a reason on its own,<sup>246</sup> few courts have taken

---

243. See Engstrom & Ho, *supra* note 56, at 841 (explaining that the SSA cannot legally disclose individual data under the Privacy Act of 1974).

244. *Id.* at 842.

245. This Article leaves aside the question of when algorithms must go through notice and comment. For a discussion of this issue, see *id.* at 845.

246. See, e.g., Jacob E. Gersen, *Legislative Rules Revisited*, 74 U. CHI. L. REV. 1705, 1719 (2007) ("Rather than asking whether a rule is legislative to answer whether notice and comment procedures should have been used, courts should simply ask whether notice and comment procedures were used. If

that approach.<sup>247</sup> Blake Emerson argues that guidance should be understood to create a “presumptively valid reason for officials to act within the policy domain it describes.”<sup>248</sup> He goes so far as to state that guidance can be entirely binding on frontline personnel, so long as there is an opportunity for the agency itself to “depart from the guidance if there are weighty reasons to do so.”<sup>249</sup> His argument parallels that of other scholars who say that agencies can treat guidance as binding on frontline personnel so long as they give regulated parties some opportunity to contest the determination and explain why they should be exempted from the guidance.<sup>250</sup>

Yet, lurking in those arguments is an assumption that guidance documents are intuitive enough to serve as reasons, whether because they are sensible on their face or because they take the form of long memoranda that include the systemic reasons for the guidance. But as discussed above, algorithms create an intuition gap when acting as reasons, making the use of algorithms distinguishable from the use of other kinds of guidance.<sup>251</sup> Because of the intuition gap, the systemic reasons for an algorithm are especially important to the algorithm’s validity as a reason. Absent at least some disclosure of systemic reasons for an algorithm, agency decisions that are based solely on the output of said algorithm should not survive arbitrary and capricious review. At a minimum, the agency should disclose what factors the algorithm considers, what testing was done, and why the use of an algorithm was important. Courts can weigh the adequacy of such justifications using the normative framework described above.

---

they were, the rule should be deemed legislative and binding if otherwise lawful. If they were not, the rule is nonlegislative. If the rule is nonlegislative, a party may challenge the validity of the rule in any subsequent enforcement proceeding; if the rule is legislative, the agency may rely on the rule in a subsequent enforcement proceeding without defending it.”). David L. Franklin has called this approach the “short cut.” David L. Franklin, *Legislative Rules, Nonlegislative Rules, and the Perils of the Short Cut*, 120 *YALE L.J.* 276, 289 (2010).

247. See Franklin, *supra* note 246, at 280 (“The federal courts themselves have never explained why they have not adopted the short cut . . .”). *But see* *Texas v. United States*, 809 F.3d 134, 175-76 (5th Cir. 2015) (holding Deferred Action for Parents of Americans and Lawful Permanent Residents (DAPA) memorandum unlawful because individual adjudicators did not retain case-by-case discretion to depart from the memorandum’s criteria for enforcement).

248. Emerson, *supra* note 223, at 2133.

249. *Id.* at 2135.

250. See, e.g., Ronald M. Levin, *Rulemaking and the Guidance Exemption*, 70 *ADMIN. L. REV.* 263, 305 (2017) (“[A]n agency should be allowed, without resorting to notice and comment, to issue a guidance document that is binding on its staff if persons affected by the document will have a fair opportunity to contest the document at a later stage in the implementation process.”).

251. See *supra* Section III.B.

CONCLUSION

This Article considered the extent to which federal agencies' use of algorithms could be reconciled with the reason-giving requirements of administrative law. While the answer depends on the algorithm and the context in which it is being used, this Article concluded that at least some algorithms should survive arbitrary and capricious review. Relatedly, this Article argued that the use of algorithms could be normatively justified despite algorithms' effects on reason-giving. In reaching these conclusions, this Article modeled a new approach to answering questions about agencies' use of algorithms: the algorithms-as-policies framework.