III. THE MERITS OF MACHINE LEARNING IN THE REGULATORY STATE

The legal analysis presented in the previous Part indicates that when federal agencies use artificial intelligence to automate regulatory and adjudicatory decisions, they will likely face little difficulty in making machine-learning practices fit within existing administrative and constitutional constraints.[275] This analysis has taken place against a backdrop of growing concern over the proliferation of machine learning and artificial intelligence, which may make our conclusions surprising, if not a bit disquieting, to some readers. After all, as we noted at the outset of this Article, technologists and sociologists alike worry

272. NLRB v. Sears, Roebuck & Co., 421 U.S. 132, 138 (1975).

273. NLRB v. Robbins Tire & Rubber Co., 437 U.S. 214, 242 (1978).

274. As discussed *infra* Section I.C, these algorithms may be embedded or tied together in overarching computer programs that translate predictions into administrative actions. Forthcoming research suggests that optimal transparency and accountability may also require that agencies publish cryptographic commitments of these programs to demonstrate the regularity of their decision making. *See* Kroll et al., *supra* note 1, at 18–21.

275. In other words, we suggest that the administrative use of machine learning will be unexceptional from a legal perspective; it will not necessitate any "systemic change to laws or legal institutions in order to preserve or rebalance established values." Calo, *supra* note 120, at 553. Interestingly, Calo describes how robotics may possess a moderate level of exceptionalism in other legal domains, such as criminal and tort law. *Id.* at 552–55.

about the consequences of algorithmic automation spreading throughout society, bringing with it the possibility that highly-skilled jobs will be taken over by machines[276] as super-intelligent computers surpass the capacities of humans.[277] These anxieties will undoubtedly grow only more fevered when artificial intelligence falls into the hands of government officials, such as when it is used in predictive policing[278] or, in what has been called "the crossing of a moral Rubicon," in target-selecting military drones.[279] If the significant governmental power wielded by administrative agencies could also be taken over by autonomous algorithms, then surely society would face a grave specter of lives and livelihoods being regulated by robots.

Notwithstanding these ominous warnings, actual technological capabilities are hardly so threatening. For much the same reason that science alone can never "make" policy decisions,[280] machine-learning algorithms need humans to specify their objective functions and construct the mathematical processes that will maximize them.[281] Although machine learning could replace or supplement many routine governmental tasks, the oversight and direction of the government will remain in human hands even in the machine-learning era. Society's most consequential regulatory decisions are not routine and therefore will almost surely prove to be unsuitable candidates for automation; these most significant regulatory policy decisions present complexities, uncertainties, and value judgments that will resist the kind of specification needed to embed them in mathematical objective functions. Machine-learning analysis will be able to assist only by informing the most significant of regulatory choices, not by determining them.[282] The more routine decisions that algorithms will be able to

---

276. *See, e.g.*, Timothy Aeppel, *What Clever Robots Mean for Jobs*, WALL ST. J. (Feb. 24, 2015, 10:30 PM), http://www.wsj.com/articles/what-clever-robots-mean-for-jobs-1424835002 [https://perma.cc/E8JJ-VB6Y]; Claire Cain Miller, *Can an Algorithm Hire Better Than a Human?*, N.Y. TIMES (June 25, 2015), http://www.nytimes.com/2015/06/26/upshot/can-an-algorithm-hire-better-than-a-human.html [https://perma.cc/N6QZ-AD55]; Zeynep Tufekci, *The Machines Are Coming*, N.Y. TIMES (Apr. 18, 2015), http://www.nytimes.com/2015/04/19/opinion/sunday/the-machines-are-coming.html [https://perma.cc/7WY5-ENV6].

277. *See, e.g.*, Peter Holley, *Apple Co-Founder on Artificial Intelligence: 'The Future Is Scary and Very Bad for People*,*'* WASH. POST (Mar. 24, 2015), https://www.washingtonpost.com/blogs/the-switch/wp/2015/03/24/apple-co-founder-on-artificial-intelligence-the-future-is-scary-and-very-bad-for-people [https://perma.cc/8ATY-TZK3].

278. *See* MAYER-SCHÖNBERGER & CUKIER, *supra* note 8, at 158–62; *see also* Barocas & Selbst, *supra* note 179, at 673; Crawford & Schultz, *supra* note 119, at 103–05; Andrew Guthrie Ferguson, *Big Data and Predictive Reasonable Suspicion*, 163 U. PA. L. REV. 327, 329–30 (2015); Ferguson, *supra* note 16 (manuscript at 3).

279. Robert H. Latiff & Patrick J. McCloskey, *With Drone Warfare, America Approaches the Robo-Rubicon*, WALL ST. J. (Mar. 14, 2013, 7:37 PM), https://www.wsj.com/articles/SB10001424127887324128504578346333246145590 [https://perma.cc/MV88-F9BQ].

280. *See* Cary Coglianese & Gary E. Marchant, *Shifting Sands: The Limits of Science in Setting Risk Standards*, 152 U. PA. L. REV. 1255, 1257–58 (2004).

281. *See supra* Section I.B.

282. Sometimes the use of benefit–cost analysis by administrative agencies has been resisted out of concern that it will substitute mechanistically for human judgment. *See, e.g.*, Lisa Heinzerling, *Regulatory Costs of Mythic Proportions*, 107 YALE L.J. 1981, 2070 (1998) (arguing that numbers "[a]t

make will surely affect people's lives and livelihoods, but in these cases machine-learning applications will be designed to meet human officials' specifications and will ultimately remain under human control. Autonomous robots will not be wielding regulatory power on their own.

For the reasons we have explained, nothing about algorithms makes them uniquely or automatically unsuitable for use by administrative agencies in terms of their ability to comport with core legal principles. Although machine-learning systems can fit quite comfortably within existing constitutional and administrative law, this does not mean that their use will always be warranted.[283] It especially does not mean that agencies should use machine learning in a haphazard or irresponsible manner. Agencies will have to act with care and, in designing and implementing their algorithms, consider potential pitfalls and areas of concern. In the following two sections, we briefly discuss how several key policy issues and good government principles apply to machine learning applications. The policy issues relate to the four legal doctrines discussed earlier: nondelegation, due process, antidiscrimination, and transparency. The other principles of good government apply more generally to the exploitation of algorithms by those in positions of power. Only by engaging in a thorough, case-by-case evaluation of such non-binding but vital considerations will agencies ensure that their use of machine learning conforms not just to the law but also to foundational principles of sound and legitimate public policy.

A. RELATED PUBLIC POLICY CONCERNS

In the previous Part, we demonstrated that, in principle, agencies' machine-learning algorithms should withstand legal challenges based on doctrines of nondelegation, due process, antidiscrimination, and transparency. Of course, an escape from judicial censure does not necessarily guarantee that the algorithms will be fully consonant with the public policy principles underlying those doctrines. Some of the policy concerns animating the legal doctrines discussed

---

worst" can "derail thoughtful discussion by offering the illusion of objective accuracy"). But even economists recognize that benefit–cost analysis can never fully determine a policy decision. *See, e.g.*, Kenneth J. Arrow et al., *Is There a Role for Benefit-Cost Analysis in Environmental, Health, and Safety Regulation?*, 272 SCI. 221, 221 (1996) (acknowledging that "benefit–cost analysis has a potentially important role to play in helping inform regulatory decision-making, although it should not be the sole basis for such decision-making"); John J. Donohue III, *Why We Should Discount the Views of Those Who Discount Discounting*, 108 YALE L.J. 1901, 1910 (1999) (noting that no one should "be a slave to such an analysis").

283. Nor will the use of machine learning always be easy. Although we have discussed how the need for specified objective functions will be more likely to produce legislation that assuages legal concerns over nondelegation, and how the same need for goal precision can mitigate concerns based on antidiscrimination and transparency principles, it may often be more difficult in practice for decision makers to agree on how to specify the objectives of administrative algorithms. As Justice Cuéllar has noted, these difficulties may arise frequently "because agreement at a high level of generality rarely translates into consensus on how to implement policies through administrative agencies." Cuéllar, *supra* note 25, at 16.

in Part II still deserve attention if agencies are to use machine learning in a responsible manner.

For instance, we noted that one of the nondelegation doctrine's motivating goals—ensuring political accountability—has prompted some scholars to advocate that the courts should give more deference when an agency head makes the key decision than when decisions are made by subordinate officials.[284] The courts have not embraced such a graduated approach, but its suggestion should prompt agency officials to consider how machine learning may affect the locus of decision making inside their agencies. When rules are intimately tied to the outputs of an algorithm, the programming of that algorithm will be a consequential task—one that presumably should not be assigned exclusively to a lower-level analyst, as traditional statistical analyses may be today. Assigning responsibility to lower-level analysts without adequate input and oversight could run the risk that higher-level officials—those who are more directly accountable to the political branches of government and to the public—will not fully understand critical details about an already ostensibly opaque rulemaking process. Perhaps worse, the lower-level analysts could make choices about an algorithm's specifications and tuning without realizing potentially far-reaching impacts of their decisions.[285] For these reasons, agencies should be mindful of who within an agency actually wields algorithm-specifying power and how well high-level officials understand the methods of regulating by robot or adjudicating by algorithm.

Our discussion of due process similarly suggests that there is no reason to view adjudicatory algorithms as uniquely threatening to fair hearings. Still, agencies will need to proceed thoughtfully and with care.[286] Agencies may need to delay the deployment of their algorithms to get a sense of how well the test-data error rates correspond with decision reversal rates—the metric underlying the existing body of relevant case law. Furthermore, agencies should increasingly seek out and engage neutral statistical experts to provide dispassionate assessments of consequential uses of algorithms. Individuals challenging an agency's deprivation of their rights or property cannot be expected to mount a sufficiently probing search for an algorithm's potential inadequacies and biases, either on their own or with the help of skilled, but mathematically naïve, counsel. Even if not legally required, agencies should still undertake the kind of probing inquiry needed to minimize possible errors and biases associated with any algorithms they use.

284. *See supra* note 140 and accompanying text.

285. In calling for increased artificial intelligence expertise in government, a recent report from Stanford notes that "insufficiently trained officials may simply take the word of industry technologists and green light a sensitive application [of artificial intelligence] that has not been adequately vetted." PETER STONE ET AL., ARTIFICIAL INTELLIGENCE AND LIFE IN 2030, at 43 (2016), https://ai100.stanford.edu/sites/default/files/ai_100_report_0901fnlb.pdf [https://perma.cc/9HPP-72XZ].

286. *See supra* Section II.B.

Agencies must also take note of the potential for their algorithms to cause a disproportionate impact on members of certain classes or groups, despite the probable constitutionality under the Fifth Amendment of the use of such algorithms.[287] Of all the policy concerns related to the legal doctrines discussed in this Article, the possibility of disproportionate impacts may be the most acute. In the popular and academic media, commentators have noted the ways in which algorithmic bias can manifest itself.[288] A series of violent interactions between police and African-Americans over the last few years has prompted public worry of pervasive discrimination by governmental authorities.[289] Agencies should thus seek ways of mitigating algorithms' disparate impact. But they should do so not just as a response to this public alarm, but also as a matter of ethical governance. If discrimination, no matter how unintentional, can be avoided, it should be. Recent advances in statistics, as noted earlier, may provide agencies with powerful mitigating tools, but agencies will still need to balance tradeoffs with forecasting accuracy.[290]

Finally, we have shown that agencies should have no problem disclosing sufficient information to meet the reason-giving and transparency requirements of the APA and FOIA.[291] Nevertheless, disclosing the bare minimum—probably just the objective functions and limited aspects of algorithms' specifications—should not be the ultimate goal of an open, forthcoming administrative state. Agencies should begin developing practices for documenting, retaining, and disclosing developmental algorithm specifications and final algorithm supplemental output to increase transparency and facilitate peer review.[292]

B. OTHER GOOD GOVERNMENT PRINCIPLES

In addition to policy concerns related to the four legal doctrines analyzed in Part II, several other good government principles should be considered by agencies when using machine learning. These principles relate to challenges associated with quantification, permissible error, lack of empathy, job losses, and privacy.

---

287. *See supra* Section II.C.

288. *See e.g.*, Barocas & Selbst, *supra* note 179; Pauline T. Kim, *Data-Driven Discrimination at Work*, 58 WM. & MARY L. REV. (forthcoming 2017); Angwin et al., *supra* note 232; Kate Crawford, Opinion, *Artificial Intelligence's White Guy Problem*, N.Y. TIMES (June 25, 2016), http://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html [https://perma.cc/RCH6-MUKF]; Noyes, *supra* note 180; Schrage, *supra* note 180.

289. *See generally* Niraj Chokshi, *How #BlackLivesMatter Came to Define a Movement*, N.Y. TIMES (Aug. 22, 2016), http://www.nytimes.com/2016/08/23/us/how-blacklivesmatter-came-to-define-a-movement.html [https://perma.cc/34W8-GNRK] (chronicling the development of the Black Lives Matter movement in response to police violence).

290. *See supra* note 229 and accompanying text.

291. *See supra* Section II.D.

292. *See supra* Section II.D. Note that OMB guidelines have stated that technical methodology subjected to peer review can be presumed to be sufficiently objective. *See* Guidelines for Ensuring and Maximizing the Quality, Objectivity, Utility, and Integrity of Information Disseminated by Federal Agencies, 67 Fed. Reg. 8,452 (Feb. 22, 2002).

First, agencies should recognize that the use of algorithms will often compel agency decision makers to engage in quantitative coding of value judgments that have typically been made qualitatively. For example, machine-learning algorithms often permit the specification of their "cost" ratios—the ratio of false positives to false negatives.[293] This is a concrete, unambiguous quantification of agencies' relative normative values of their errors. For agencies not accustomed to making moral valuations through any kind of formal process, let alone one that assigns them numbers, machine-learning algorithms will necessitate addressing questions of organizational and democratic decision making. Who should have the power to transform qualitative moral judgments into a cost ratio? Should agencies involve the public or at least make interested parties aware of a transformation from qualitative to a quantitative assessment of errors and algorithmic tradeoffs? Is such a transformation even possible, or are human deliberations over morality too nuanced to be reduced to a matrix of error rates?[294] Agencies will have to confront not only the challenges of making accurate quantifications of error rates but also contend with a host of relevant normative questions about quantification as they proceed into the machine-learning era.

Second, agencies will have to decide what constitutes "acceptable" algorithmic error rates. This will go beyond the due process-related assessment of whether courts will consider algorithm error rates acceptable. Instead, it requires agencies to ask how large gains in accuracy must be to offset other difficulties in implementing algorithmic systems. Such issues would likely arise in the context of benefit–cost analyses that agencies conduct prior to taking important actions. For example, how much would, say, a five percent increase in accuracy when predicting hazardous waste pipeline incidents save an agency in inspection costs? Would these savings be worth the necessary investments in human capital and data infrastructure? And, beyond the internal difficulties faced by agencies, would the public demand particularly large increases in accuracy to compensate for the novelty, and potentially alarming nature, of robotic systems of regulation or adjudication?[295] Is a marginal improvement over the status quo acceptable, or should agencies strive for something more?

293. *See supra* note 172.

294. The difficulties of programming morality into artificially intelligent systems has been discussed previously, such as in the context of the choice of autonomous cars about which humans to harm when faced with scenarios inevitably resulting in some human casualties. *See, e.g.*, John Markoff, *Should Your Driverless Car Hit a Pedestrian to Save Your Life?*, N.Y. TIMES (June 23, 2016), http://www.nytimes.com/2016/06/24/technology/should-your-driverless-car-hit-a-pedestrian-to-save-your-life.html [https://perma.cc/PL2N-R7GR].

295. Governmental use of machine learning could create a "Tesla effect," of sorts. The Tesla effect refers to the principle that, even though driverless cars may be statistically far safer than manned ones, widespread alarm results when even a few individuals become victims of accidents involving a Tesla operating autonomously. *Cf.* Larry Greenemeier, *Deadly Tesla Crash Exposes Confusion over Automated Driving*, SCI. AM. (July 8, 2016), http://www.scientificamerican.com/article/deadly-tesla-crash-exposes-confusion-over-automated-driving [https://perma.cc/72C4-T33P] (discussing both self-

Third, we should consider what happens to human contact in a world of robotic regulators and algorithmic adjudicators. Citizens expect their government to be competent and efficient, but they also tend to view governmental institutions as more legitimate when these institutions operate with understanding and empathy.[296] The idea of the government reducing individuals to data points that are then fed into an algorithm will seem disconcertingly impersonal—even if ultimately more accurate and efficient. Administrative officials that use algorithms should seek to listen to interested members of the public as they design systems for rulemaking by robot or adjudicating by algorithm. These officials should also encourage participation through interactive methods that treat beneficiaries, targets of regulation, and all other affected parties with respect and dignity.[297]

Fourth, government should not turn a blind eye to the possibility that widespread use of algorithms will put many government workers out of their jobs.[298] The extent to which agencies' workforces will shrink as a result of adopting machine learning is unclear; as we have noted, human input will still be necessary at many steps of regulatory and adjudicatory processes, even when automated.[299] But there can be little doubt that the structures of agencies' workforces will change if day-to-day operations begin to take drastically different forms due to automated decision making. With over 1.8 million full-time

---

driving's safety features and the collective apprehension after a fatal accident involving a self-driving car).

296. *See, e.g.*, Tom R. Tyler, *Procedural Justice, Legitimacy, and the Effective Rule of Law*, 30 CRIME & JUSTICE 283 (2003).

297. *See* CARY COGLIANESE, LISTENING, LEARNING, LEADING: A FRAMEWORK FOR REGULATORY EXCELLENCE 9 (2015), https://www.law.upenn.edu/live/files/4946-pprfinalconvenersreport.pdf [https://perma.cc/F89X-SGHB] (describing "empathic engagement" as a core attribute of regulatory excellence); Cary Coglianese, *Regulatory Excellence as "People Excellence,"* REGBLOG (Oct. 23, 2015), http://www.regblog.org/2015/10/23/coglianese-people-excellence/ [https://perma.cc/7AND-FB8K] ("[R]egulatory excellence demands the consistent achievement of three fundamental attributes: *utmost integrity, empathic engagement,* and *stellar competence*." (emphasis in original)). It may even be worthwhile for agencies to consider ways of developing online, anthropomorphic representations as part of a human–computer interface to encourage more empathic and emotionally positive interactions between citizens and their increasingly robotic regulators. *Cf.* Kate Darling, *"Who's Johnny?": Anthropomorphic Framing in Human-Robot Interaction, Integration, and Policy*, *in* ROBOT ETHICS 2.0 (forthcoming 2017), http://ssrn.com/abstract=2588669 [https://perma.cc/66PQ-6U85] (recognizing concerns around framing robotic technology in human terms but noting the benefits of anthropomorphizing robots); Will Davies, *Robot Amelia—A Glimpse of the Future for Local Government*, GUARDIAN (July 4, 2016, 2:10 AM), https://www.theguardian.com/public-leaders-network/2016/jul/04/robot-amelia-future-local-government-enfield-council [https://perma.cc/EZG3-GXCF] (describing how a London borough has begun responding to citizen requests, such as for permits, using a voice response system that employs natural language processing to interpret emotions in citizens' voices and respond appropriately and empathetically); Adriana Hamacher et al., Believing in BERT: Using Expressive Communication to Enhance Trust and Counteract Operational Error in Physical Human-Robot Interaction, Presented at the IEEE International Symposium on Robot and Human Interactive Communication (2016), http://arxiv.org/abs/1605.08817 [https://perma.cc/PV72-UZFM] (describing how robots with affective interaction styles recover more of humans' trust after they make errors than robots that are more efficient yet impersonal).

298. *See supra* note 9 and accompanying text.

299. *See supra* Section I.C.

employees in the executive branch of the federal government,[300] and many more working for government contractors, large-scale automation efforts may be appropriately combined with proactive job training for workers who might seek positions in the future in algorithm-support responsibilities.

Finally, although throughout this Article we have self-consciously eschewed a discussion of privacy concerns, for reasons noted earlier,[301] agencies still must never overlook these concerns. Administrative agencies are poised to collect more data on individuals with each passing day. Agencies must properly and securely store these data to minimize threats to privacy intrusions, especially when many administrative applications of machine learning will require inter-agency sharing through the cloud.[302] Agencies may also increasingly seek individuals' data from sources outside of the United States, which will require careful consideration of jurisdictional questions.[303]

C. A PATH FORWARD

The issues we have raised in the preceding two sections should be addressed as part of agencies' case-by-case assessments and benefit–cost analyses of specific applications of machine learning. Agency and cross-agency decisions about when and how to implement machine learning will benefit from clear guidance on how to assess machine learning's merits.[304] The Administrative Conference of the United States, the National Academy of Public Administration, or the National Academy of Sciences might be able to help in developing guidelines. When agencies ultimately apply such guidelines and conduct benefit–cost analysis of specific machine learning applications, we can expect administrative algorithms will prove extremely promising in some instances and less advantageous in others.[305] Agencies should be suitably discerning and pursue

300. U.S. Office of Pers. Mgmt., SIZING UP THE EXECUTIVE BRANCH: FISCAL YEAR 2015, at 5 (2016), https://www.opm.gov/policy-data-oversight/data-analysis-documentation/federal-employment-reports/reports-publications/sizing-up-the-executive-branch-2015.pdf [https://perma.cc/4FGX-2B4H].

301. *See supra* note 119 and accompanying text.

302. *Cf.* Paul M. Schwartz, *Information Privacy in the Cloud*, 161 U. PA. L. REV. 1623, 1661–62 (2013) (considering United States and European Union definitions of "personal information" and suggesting regulatory reforms for cloud storage to ensure "strong and effective protections for information privacy").

303. *Cf.* Andrew Keane Woods, *Against Data Exceptionalism*, 68 STAN. L. REV. 729, 731 (2016) (discussing competing conceptions of the territoriality of personal data).

304. This kind of ethical oversight has also been called for in recent considerations of a future medical profession driven substantially by machine learning. *See* Alison M. Darcy et al., Opinion, *Machine Learning and the Profession of Medicine*, 315 J. AM. MED. ASS'N. 551, 551 (2016) ("The profession of medicine has a tremendous opportunity and an obligation to oversee the application of this [machine-learning] technology to patient care.").

305. There may also be instances in which pursuing machine learning is advantageous and worthwhile only when human intuition can be incorporated into algorithms. *Cf.* Jens Jakob W. H. Sørensen et al., *Exploring the Quantum Speed Limit with Computer Games*, 532 NATURE 210, 213 (2016) (describing how addressing key technical issues facing quantum computing may be facilitated by incorporating the results of human intuition, gleaned from computer games simulating atomic movement, into machine-learning algorithms).

machine learning applications when they can lead to meaningful improvements in procedural and substantive outcomes.

Recognition of the need for careful consideration and reasonable safeguards when agencies use algorithms should not lead to any presumption against using machine learning entirely. These techniques are becoming ubiquitous in private industry for good reason; their ability to make accurate predictions of complex phenomena can render decision making vastly more effective and efficient, and it would be wise for administrative agencies also to seek these benefits.[306] Still, accountable managers must carefully oversee their use of algorithms, even in the private sector,[307] and they must take possible unintended consequences into account.[308] The same can be said of all innovative practices. The need to weigh the pros and cons of algorithms is in no way qualitatively different than the weighing of benefits and costs needed to inform other administrative choices,[309] including those prompted by other technological advancements.[310] Thoughtful implementation is always advisable for the adoption of any new administrative technology or process.

Deciding how and when to use machine learning may not come easily, which may make agency efforts to facilitate public participation in decisions about the use of machine learning even more important. It would certainly be better to take additional time to engage in robust public consultation and make thoughtful decisions about machine learning's use than to dismiss algorithms out of hand over exaggerated fears about unleashing artificial intelligence "demons."[311] Technological and analytical advances are continually invented and adopted because, despite their potential limitations, they offer the possibility to transform society for the better. The implementation of machine-learning algo-

---

306. Coglianese, *supra* note 22.

307. *See* Michael Luca, Jon Kleinberg & Sendhil Mullainathan, *Algorithms Need Managers, Too*, 94 HARV. BUS. REV. 96 (Jan.–Feb. 2016), https://hbr.org/2016/01/algorithms-need-managers-too [https://perma.cc/7YAB-MK4S].

308. *Cf*. Gökçe Sargut & Rita McGrath, *Learning to Live with Complexity*, HARV. BUS. REV. (Sept. 2011), https://hbr.org/2011/09/learning-to-live-with-complexity [https://perma.cc/94HV-65PW] (describing the management techniques required to oversee complex, data-driven systems); Latanya Sweeney, *Discrimination in Online Ad Delivery*, 56 COMMC'NS ACM 44, 53 (2013) (describing Google's need to contend with differential delivery of advertisements for arrest records when individuals search for names typically associated with different races). Algorithms applied in administrative contexts will face the same possible risks of bounded cognition that can accompany well-accepted, performance-based regulatory standards because objective functions will be defined in terms of desired outcomes. *See* Cary Coglianese, *Performance-Based Regulation: Concepts and Challenges*, *in* COMPARATIVE LAW AND REGULATION: UNDERSTANDING THE GLOBAL REGULATORY PROCESS 403 (Francesca Bignami & David Zaring eds., 2016).

309. *Cf.* Charles E. Lindblom, *The Science of "Muddling Through,"* 19 PUB. ADMIN. REV. 79 (1959) (discussing the complexity of the mosaic nature of traditional administrative decision making).

310. The Internet, for example, has vastly expanded the ways in which administrative agencies can communicate with the public in rulemaking, but taking advantage of these opportunities still requires careful consideration of factors such as how to make information equally accessible to all members of the public. *See* Cary Coglianese, *Enhancing Public Access to Online Rulemaking Information*, 2 MICH. J. ENVTL. & ADMIN. L. 1, 39–40 (2012).

311. *Cf*. Gibbs, *supra* note 11.

rithms by administrative agencies certainly possesses such potential. With proper forethought and care guided by the kind of analysis we have provided throughout this Article, an administrative state on the cutting edge of statistical innovation can be legally and responsibly realized. Moving in this direction can deliver marked improvements in overall well-being, especially as the government faces new challenges in overseeing a private sector that has already entered the machine-learning era.

312. Coglianese, *supra* note 22.

313. *See* Brenden M. Lake et al., *Human-Level Concept Learning Through Probabilistic Program Induction*, 350 SCIENCE 1332 (2015) (suggesting that the brain–computer analogy could profitably guide research); Gary Marcus, *Face It, Your Brain Is a Computer*, N.Y. TIMES (June 27, 2015), http://www.