

Administrative Law in the Automated State

Cary Coglianese

In the future, administrative agencies will rely increasingly on digital automation powered by machine learning algorithms. Can U.S. administrative law accommodate such a future? Not only might a highly automated state readily meet long-standing administrative law principles, but the responsible use of machine learning algorithms might perform even better than the status quo in terms of fulfilling administrative law's core values of expert decision-making and democratic accountability. Algorithmic governance clearly promises more accurate, data-driven decisions. Moreover, due to their mathematical properties, algorithms might well prove to be more faithful agents of democratic institutions. Yet even if an automated state were smarter and more accountable, it might risk being less empathic. Although the degree of empathy in existing human-driven bureaucracies should not be overstated, a large-scale shift to government by algorithm will pose a new challenge for administrative law: ensuring that an automated state is also an empathic one.

Because the future knows no bounds, the future of administrative law is vast indeed. In the near term, administrative law in the United States will undoubtedly center around how the U.S. Supreme Court decides cases raising core administrative law issues such as the nondelegation doctrine and judicial deference to agencies' statutory interpretation. But over the longer term, new issues will confront the field of administrative law as new changes occur in government and in society. One major change on the horizon will be an increasingly automated administrative state in which many governmental tasks will be carried out by digital systems, especially those powered by machine learning algorithms.

Administrative agencies today undertake a range of activities – granting licenses, issuing payments, adjudicating claims, and setting rules – each of which traditionally has been executed by government officials. But it is neither difficult nor unrealistic to imagine a future in which members of the public, when they interact with government, increasingly find themselves interacting predominantly with digital systems rather than human officials. Even today, the traditional administrative tasks for which human beings have long been responsible are increasingly augmented by computer systems. Few people in the United States today think

twice about using government websites to apply for unemployment benefits, register complaints, or file paperwork, rather than visiting or telephoning government offices. The federal government has even created an online portal – USA.gov – that provides its users with easy access to the panoply of resources and digital application processes now available to the public via an extensive network of state and federal government websites.

The transition to this online interaction with government over the last quarter-century portends what will likely be a deeper and wider technological transformation of governmental processes over the next quarter-century. Moving beyond the digitization of front-end communication with government, the future will likely feature the more extensive automation of back-end decision-making, which today still often remains firmly in the discretion of human officials. But we are perhaps only a few decades away from an administrative state that will operate on the basis of automated systems built with machine learning algorithms, much like important aspects of the private sector increasingly will. This will lead to an administrative state characterized by what I have elsewhere called *algorithmic adjudication* and *robotic rulemaking*.¹ Instead of having human officials make discretionary decisions, such as judgments about whether individual claimants qualify for disability benefits, agencies will be able to rely on automated systems to make these decisions. Claims-processing systems could be designed, for example, to import automatically a vast array of data from electronic medical records and then use an artificial intelligence system to process these data and determine whether claimants meet a specified probability threshold to qualify for benefits.²

If many of the tasks that government currently completes through decision-making by human officials come to be performed entirely by automated decision tools and computer systems, how will administrative law respond to this transformation to an automated state? How *should* it?

Most existing administrative law principles can already accommodate the widespread adoption of automation throughout the administrative state. Not only have agencies already long relied on a variety of physical machines that exhibit automaticity, but an automated state – or at least a *responsible* automated state – could be thought of as the culmination of administrative law’s basic vision of government that relies on neutral public administration of legislatively delegated authority. Administrative law will not need to be transformed entirely to operate in an era of increasing automation because that automation, when responsibly implemented, will advance the democratic principles and good governance values that have long underlay administrative law.

Nevertheless, even within an otherwise responsible automated state, there will come to be an important ingredient of good governance that increasingly could turn out to be missing: human empathy. Even bureaucracies comprising human officials can be cold and sterile, but an era of extreme automation could present a

state of crisis in human care – or, more precisely, a crisis in the lack of such care. In an increasingly automated state, administrative law will need to find ways to encourage agencies to ensure that members of the public will continue to have opportunities to engage with humans, express their voices, and receive acknowledgment of their predicaments. The automated state will, in short, also need to be an empathic state.

The information technology revolution that launched several decades ago shows few signs of abating. Technologists today are both revealing and reaching new frontiers with the use of advanced algorithmic technologies variously referred to as artificial intelligence, machine learning, and predictive analytics. These terms – sometimes used interchangeably – encompass a broad range of tools that permit the rapid processing of large volumes of data that can yield highly accurate forecasts and thereby facilitate the automation of many distinct tasks. In the private sector, algorithmic innovations are allowing the automation of a wide range of functions previously handled by trained humans, such as the reading of chest X-rays, the operation of automobiles, and the granting of loans by financial institutions.

Public administrators have taken notice of these algorithmic advances in the private sector. Some advances in the business world even have direct parallels to governmental tasks. Companies such as eBay and PayPal, for example, have developed their own highly successful automated online dispute resolution tools to resolve complaints without the direct involvement of human employees.³ Overall, government officials see in modern data analytics the possibility of building systems that could automate a variety of governmental tasks, all with the potential to deliver increased administrative efficiency, speed, consistency, and accuracy.

The vision of an automated administrative state might best be exemplified today by developments in the Republic of Estonia, a small Baltic country that has thoroughly embraced digital government as a mark of distinction. The country's e-Estonia project has transformed the nation's administration by digitizing and securely storing vast amounts of information about individuals, from their medical records to their employment information to their financial statements.⁴ That information is cross-linked through a digital infrastructure called X-Road, so that a person's records can be accessed instantly by any entity that needs them, subject to limits intended to prevent wrongdoing. This widespread digitization has facilitated the automation of a range of government services: individuals can easily vote, apply for a loan, file their taxes, and complete other administrative tasks without ever needing to interact with a human official, simply by transferring their digital information to complete forms and submit requests. By automating many of its bureaucratic processes, Estonia has saved an estimated 2 percent of its

GDP each year. The country is even exploring the use of an automated “judge” to resolve small claims disputes.⁵

Other countries such as Denmark and South Korea are also leading the world in the adoption of so-called e-government tools.⁶ The United States may not have yet achieved quite the same level of implementation of automated government, but it is certainly not far behind. Federal, state, and local agencies throughout the United States have not only embraced web-based applications – such as those compiled on the USA.gov website – but have begun to deploy the use of machine learning algorithms to automate a range of administrative decision-making processes. In most of these cases, human officials remain involved to some extent, but a significant amount of administrative work in the United States is increasingly conducted through digital systems.

Automation helps federal, state, and local governments navigate challenging resource-allocation decisions in the management of public programs. Several states have implemented algorithmic tools to help make decisions about the award of Medicaid and other social benefits, seeking to speed up and improve the consistency of claims processing.⁷ Similarly, the federal Social Security Administration uses automated tools to help support human appeals judges’ efforts to provide quality oversight of an agency adjudicatory process that handles as many as 2.5 million disability benefits claims each year.⁸

Municipalities rely on automated systems when deciding where to send health and building inspectors.⁹ Some local authorities use such systems when making choices about where and when to deploy social workers to follow up on allegations of child abuse and neglect.¹⁰ Federal agencies, meanwhile, have used algorithmic systems to analyze consumer complaints, process reports of workplace injuries, and evaluate public comments on proposed rules.¹¹

Criminal law enforcement agencies throughout the United States also rely on various automated tools. They have embraced tools that automate deployment of officer patrols based on predictions of locations in cities where crime is most likely to occur.¹² Many law enforcement agencies have also widely used automated facial recognition tools to facilitate suspect identification or for security screenings.¹³

Regulatory agencies similarly have deployed automated tools for targeting auditing and enforcement resources. States have employed data analytics to detect fraud and errors in their unemployment insurance programs.¹⁴ The federal Securities and Exchange Commission and the Internal Revenue Service have adopted algorithmic tools to help detect fraudulent behavior and other wrongdoing.¹⁵

In these and other ways, public authorities across the United States have already made considerable strides toward an increasingly automated government. Over the next several decades, governmental use of automation driven by artificial intelligence tools will surely spread still further and is likely to lead to the transformation of or phasing out of many jobs currently performed by govern-

ment employees.¹⁶ The future state that administrative law will govern will be one of increasingly automated administration.

Can administrative law accommodate an automated state? At first glance, the prospect of an automated state might seem to demand a fundamental re-writing of administrative law. After all, administrative law developed to constrain the discretion of human officials, to keep their work within the bounds of the law, and to prevent the kinds of principal-agent problems that can arise in the relationships between human decision-makers. Moreover, one of administrative law's primary tenets – that governmental processes should be transparent and susceptible to reason-giving – would seem to stand as a barrier to the deployment of the very machine learning algorithms that are driving the emerging trends in automation.¹⁷ That is because machine learning algorithms – sometimes referred to as “black-box” algorithms – have properties that can make them opaque and hard to explain. Unlike traditional statistical algorithms, in which variables are selected by humans and resulting coefficients can be pointed to as explaining specified amounts of variation in a dependent variable, learning algorithms effectively discover their own patterns in the data and do not generate results that associate explanatory power to specific variables. Data scientists can certainly understand and explain the goals and general properties of machine learning algorithms, but overall these algorithms have a degree of autonomy – hence their “learning” moniker – that can make it more difficult to explain precisely why they reach any specific forecast that they do. They do not usually provide any basis for the kind of causal statements often used to justify administrative decisions (such as “X is justified because it causes Y”).

As a result, transparency concerns are reasonable when considering a future of an automated state based on machine learning systems. But on even a modest degree of additional reflection, these concerns would appear neither to act as any intrinsic barrier to the reliance on machine learning automation nor necessarily to demand any fundamental transformation of U.S. administrative law to accommodate an automated state. Administrative law has never demanded anything close to absolute transparency nor required meticulous or exhaustively detailed reasoning, even under the arbitrary and capricious standard of Section 706 of the Administrative Procedure Act.¹⁸ Administrative agencies that rely on machine learning systems should be able to satisfy any reason-giving obligations under existing legal principles by explaining in general terms how the algorithm was designed to work and demonstrating that it has been validated to work as designed by comparing its results to those generated by the status quo process. An adequate explanation could involve merely describing the type of algorithm used, disclosing the objective the algorithm was established to meet, and showing how the algorithm processed a certain type of data to produce results that were shown to meet the algorithm's defined objective as well as or better than current processes.

Such an explanation would, in effect, mirror the kinds of explanations that administrators currently offer when they rely on physical rather than digital machines. For example, in justifying the imposition of an administrative penalty on a food processor for failing to store perishable food at a cool temperature, an administrator need not be able to explain exactly how a thermometer works, just that it reports temperatures accurately. Courts have long treated instrument validation for physical machines as a sufficient basis for agency actions grounded on such instruments. Moreover, they have typically deferred to administrators' expertise in cases in which government officials have relied on complex instruments or mathematical analyses. In fact, the U.S. Supreme Court in *Baltimore Gas & Electric Co. v. Natural Resources Defense Council* called upon courts to be their "most deferential" when an administrative agency is "making predictions, within its area of special expertise, at the frontiers of science."¹⁹ More recently, the Supreme Court noted in *Marsh v. Oregon Natural Resource Council* that whenever an agency decision "requires a high degree of technical expertise, we must defer to 'the informed discretion of the responsible agencies.'"²⁰ Lower courts have followed these instructions and have upheld agencies' reliance on complex (even if not machine learning) algorithms in various contexts.

It is difficult to see the Supreme Court gaining any more confidence in judges' ability to provide independent technological assessments when technologies and statistical techniques grow still more complex in an era of machine learning. Unless the Court should gain a new source of such confidence and abandon the postures it took in *Baltimore Gas & Electric* and *Marsh*, nothing in administrative law's reason-giving requirements would seem to serve as any insuperable barrier to administrative agencies' more extensive reliance on systems based on machine learning or other advanced predictive techniques, even if they are properly characterized today as black-box algorithms. That portrayal of machine learning algorithms as a black box also appears likely to grow less apt in the coming decades, as data scientists are currently working extensively to develop advanced techniques that can better explain the outputs such complex algorithms generate.²¹ Advances in "explainable" artificial intelligence techniques likely will only make automation still more compatible with long-standing administrative law values.

Of course, all of this is not to say that agencies will or should always receive deference for how they design or operate their systems. Under the standard articulated in *Motor Vehicle Manufacturers Association v. State Farm Insurance Co.*, agencies will still need to provide basic information about the purposes behind their automated systems and how they generally operate.²² They will need to show that they have carefully considered key design options. And they will likely need to demonstrate through accepted auditing and validation efforts that these systems do operate to produce results as intended.²³ But all this is to say that it will almost certainly be possible for agencies to provide the necessary information to justify the

outcomes that their systems produce. In other words, long-standing administrative law principles seem ready and fit for an automated age.

In important respects, a shift to automated administration could even be said to represent something of an apotheosis of the principles behind administrative law. Much of administrative law has been focused on the potential problems created by the discretion that human officials exercise under delegated authority. By automating administration, those problems can be mitigated, and the control of human discretion may be enhanced by the literal hardwiring of certain governmental tasks.

Automation can advance two major themes that have long characterized much of U.S. administrative law: one theme centers on keeping the exercise of administrative authority democratically accountable, while the other seeks to ensure that such authority is based on sound expert judgment. The reason-giving thrust behind the Administrative Procedure Act's arbitrary and capricious standard, for example, reflects both of these themes. Reasoned decision-making provides a basis for helping ensure that agencies both remain faithful to their democratic mandates and base their decisions on sound evidence and analysis. Likewise, the institutionalized regimen of White House review of prospective regulations both facilitates greater accountability to a democratically elected president and promotes expert agency decision-making through the benefit-cost analysis that it calls on agencies to conduct.²⁴

In the same vein, in approving judicial deference to agencies' statutory interpretations, it is little accident that the Supreme Court's widely cited decision in *Chevron v. Natural Resources Defense Council* stressed both reasons of democratic accountability and substantive expertise.²⁵ It highlighted how agencies are situated within a "political branch of the Government" as well as how they simultaneously possess "great expertise" – and thus are better suited than courts to make judgments about the meaning of ambiguous statutory terms.²⁶ Although the future of the *Chevron* doctrine itself appears uncertain at best, the Court's underlying emphasis on accountability and expertise is unlikely to disappear, as they are inherent qualities of administrative governance.

Both qualities can be enhanced by machine learning and automation. It is perhaps most obvious that automation can contribute to the goal of expert administration. When automated systems improve the accuracy of agency decision-making – which is what makes machine learning and other data analytic techniques look so promising – this will necessarily promote administrative law's goal of enhancing agency expertise. Artificial intelligence promises to deliver the state of the art when it comes to expert governing. When the Veterans Administration (VA), for example, recently opted to rely on an automated algorithmic system to predict which veterans are at a higher risk of suicide (and thus in need of more

urgent care), it did so because this analytic system was smarter than even experienced psychiatrists.²⁷ “The fact is, we can’t rely on trained medical experts to identify people who are truly at high risk [because they are] no good at it,” noted one VA psychiatrist.²⁸

Likewise, when it comes to administrative law’s other main goal – democratic accountability – automated systems can also advance the ball. The democratic advantages of automation may seem counterintuitive at first: machine-based governance would hardly seem consistent with a Lincolnesque notion of government by “the people.” But the reality is that automated systems themselves still demand people who can design, test, and audit such systems. As long as these human designers and overseers operate systems in a manner consistent with the parameters set out for an agency in its governing statute, automated systems themselves can prevent the kind of slippage and shirking that can occur when agencies must rely on thousands of human officials to carry out major national programs and policies. Even when it comes to making new rules under authority delegated to it by Congress, agencies could very well find that automation promotes democratic accountability rather than detracts from it. Some level of accountability will be demanded by the properties of machine learning algorithms themselves. To function, these algorithms depend not merely on an “intelligible principle” to guide them; they need a principle that can be precisely specified in mathematical terms.²⁹ In this way, automation could very well drive the demand for still greater specification and clarity in statutes about the goals of administration, more than even any potential judicial reinvigoration of the nondelegation doctrine might produce.

Although oversight of the design and development of automated systems will remain important to ensure that they are created in accord with democratically affirmed values, once operating, they should pose far fewer opportunities for the kinds of problems, such as capture and corruption, that administrative law has long sought to prevent. Unlike human beings, who might pursue their own narrow interests instead of those of the broader public, algorithms will be programmed to optimize the objectives defined by their designers. As long as those designers are accountable to the public, and as long as the system objectives are defined in non-self-interested ways that comport with relevant legislation, then the algorithms themselves pose no risk of capture and corruption. In an important sense, they will be more accountable in their execution than even human officials can be when it comes to implementing law.

This is not to suggest that automated systems will amount to a panacea nor that their responsible development and use will be easy. They can certainly be used in legally and morally problematic ways. Furthermore, their use by agencies will still be subject to constraints beyond administrative law – for instance, legal constraints under the First Amendment or the Equal Protection Clause – that apply to all gov-

ernmental actions. In fact, equality concerns raised by the potential for algorithmic bias may well become the most salient legal issue that automated systems will confront in the coming years. Bias obviously exists with human decision-making, but it also is a concern with machine learning algorithms, especially when the underlying data used to train these algorithms already contain (human-created) biases. Nevertheless, absent an independent showing of animus, automated systems based on machine learning algorithms may well withstand scrutiny under equal protection doctrine, at least if that doctrine does not change much over time.³⁰

Governmental reliance on machine learning algorithms would be able to avoid actionable conduct under equal protection analysis even if an administrator elected to use data that included variables on race, gender, or other protected classifications. As long as the objective the algorithm is programmed to achieve is not stated in terms of such protected classifications, it will be hard, if not impossible, to show that the algorithm has used any class-based variables as a determinative basis for any particular outcome. The outcomes these algorithms generate derive from effectively autonomous mathematical processes that discern patterns among variables and relationships between different variables. Presumably, machine learning algorithms will seldom if ever support the kind of clear and categorical determinations based on class-related variables that the Supreme Court has rejected, where race or other protected classes have been given an explicit and even dispositive weight in governmental decisions.³¹ Even when used with data on class variables, the use of machine learning algorithms might well lead to better outcomes for members of a protected class overall.³²

Moreover, with greater reliance on algorithm-based automated systems, governments will have a new ability to *reduce* undesired biases by making mathematical adjustments to their algorithms, sometimes without much loss in accuracy.³³ Such an ability will surely make it easier to tamp out biases than it currently is to eliminate humans' implicit biases. In an automated state of the future, government may find itself less prone to charges of undue discrimination.

For these reasons, it would appear that long-standing principles of administrative law, and even constitutional law, will likely continue to operate in an automated state, encouraging agencies to act responsibly by both preserving democratic accountability and making smarter, fairer decisions. This is not to say that existing principles will remain unchanged. No one should expect that any area of the law will stay static over the long term. Given that some scholars and observers have already come to look critically upon governmental uses of algorithms, perhaps shifting public attitudes will lead to new, potentially more demanding administrative law principles specifically targeting the automated features of the future administrative state.³⁴

While we should have little doubt that norms and best practices will indeed solidify around how government officials ought to use automated systems – much

as they have developed over the years for the use of other analytic tools, such as benefit-cost analysis – it is far from clear that the fundamentals of administrative law will change dramatically in an era of algorithmic governance.³⁵ Judges, after all, will confront many of the same difficulties scrutinizing machine learning algorithms as they have confronted in the past with respect to other statistical and technical aspects of administration, which may lead to continued judicial deference as exemplified in *Baltimore Gas & Electric*.³⁶ In addition, rather than public attitudes turning against governmental use of algorithmic tools, it may just as easily be expected that public expectations will be shaped by widespread acceptance of artificial intelligence in other facets of life, perhaps even leading to affirmative demands that governments use algorithmic tools rather than continuing to rely on slower or less reliable processes. Cautious about ossifying algorithmic governance, judges and administrative law scholars might well resist the urge to impose new doctrinal hurdles on automation.³⁷ They may also conclude, as would be reasonable, that existing doctrine contains what is needed to ensure that government agencies use automated systems responsibly.

As a result, if government agencies wish to expand the responsible use of properly trained, audited, and validated automated systems that are sufficiently aligned with legislative mandates and improve agencies' ability to perform key tasks, it seems they will hardly need any transformation of traditional administrative law principles to accommodate these innovations. Nor will administrative law need to adapt much, if at all, to ensure that kind of responsible use of algorithmic governance. Overall, an automated state could conceivably do a better job than ever before of fulfilling the vision of good governance that has long animated administrative law.

Still, even if the prevailing principles of administrative law can deal adequately with public sector use of machine learning algorithms, something important could easily end up getting lost in an automated state. Such an administrative government might be smarter, more democratically accountable, and even more fair. But it could also lack feeling, even more than sterile bureaucratic processes do today. Interactions with government through smartphones and automated chats may be fine for making campground reservations at national parks or even for filing taxes. But they run the risk of leaving out an important ingredient of good governance – namely, empathy – in those circumstances in which government must make highly consequential decisions affecting the well-being of individuals. In such circumstances, empathy demands that administrative agencies provide opportunities for human interaction and for listening and expressions of concern. An important challenge for administrative law in the decades to come will be to find ways to encourage an automated state that is also an empathic state.

A desire for empathy, of course, need not impede the development of automation.³⁸ If government manages the transition to an automated state well, it is possible that automation can enhance the government's ability to provide empathy to members of the public, but only if government officials are sufficiently attentive to the need to do so. This need will become even greater as the overall economy moves toward greater reliance on artificial intelligence and other automated systems. Society will need to value and find new ways to fulfill those tasks involving empathy that humans are good at fulfilling. The goal should be, as technologist Kai-Fu Lee has noted, to ensure that, "while AI handles the routine optimization tasks, human beings . . . bring the personal, creative, and compassionate touch."³⁹

Already, public administration experts recognize that this is one of the great potential advantages of moving to an automated state. It can free up government workers from drudgery and backlogs of files to process, while leaving them more time and opportunities to connect with those affected by agency decisions.⁴⁰ A recent report jointly issued by the Partnership for Public Service and the IBM Center for Business and Government explains the importance of this shift in what government employees do:

Many observers who envision greater use of AI in government picture more face-to-face interactions between agency employees and customers, and additional opportunities for more personalized customer services. The shift toward employees engaging more with agency customers is expected to be one of several possible effects of automating administrative tasks. Relieved of burdensome paperwork, immigration officers could spend more time interacting with visa applicants or following up on individual immigration cases. Scientists could allot more of their day to working with research study participants. And grants managers could take more time to learn about and support individual grantees. On average, federal employees now spend only 2 percent of their time communicating with customers and other people outside their agencies, or less than one hour in a workweek, according to one study. At the same time, citizens want government to do better. The experiences customers have with companies is driving demand for personalized government services. In a survey of more than 6,000 people from six countries, including the United States, 44 percent of respondents identified personalized government services as a priority.⁴¹

Not only does a substantial portion of the public already recognize the need for empathic, personalized engagement opportunities with government, but as private sector organizations invest more in personalized services, this will only heighten and broaden expectations for similar empathy from government. We already know from extensive research on procedural justice that the way that government treats members of the public affects their sense of legitimacy in the outcomes they receive.⁴² To build public trust in an automated state, government authorities will need to ensure that members of the public still feel a human con-

nection. As political philosopher Amanda Greene has put it, “government must be seen to be sincerely caring about each person’s welfare.”⁴³

Can administrative law help encourage empathic administrative processes? Some might say that this is already a purpose underlying the procedural due process principles that make up administrative law. *Goldberg v. Kelly*, after all, guarantees certain recipients of government benefits the right to an oral hearing before a neutral decision-maker prior to the termination of their benefits, a right that does afford at least an opportunity for affected individuals to engage with a theoretically empathic administrative judge.⁴⁴ But the now-canonical test of procedural due process reflected in *Mathews v. Eldridge* is almost entirely devoid of attention to the role of listening, caring, and concern in government’s interactions with members of the public.⁴⁵ *Mathews* defines procedural due process in terms of a balance of three factors: 1) the affected private interests; 2) the potential for reducing decision-making error; and 3) the government’s interests concerning fiscal and administrative burdens. Machine learning automation would seem to pass muster quite easily under the *Mathews* balancing test. The first factor – the private interests at stake – will be external to machine learning, but machine learning systems would seem always to fare well under the second and third factors. Their great promise is that they can reduce errors and lower administrative costs.

This is where existing principles of administrative law will fall short in an automated state and where the need for greater vision will be needed. Hearing rights and the need for reasons are about more than just achieving accurate outcomes, which is what the *Mathews* framework implies. On the contrary, hearings and reasoning might not be all that good at achieving accurate outcomes, at least not as consistently as automated systems. A 2011 study showed that, among the fifteen most active administrative judges in one office of the Social Security Administration, “the judge grant rates . . . ranged . . . from less than 10 percent being granted to over 90 percent.”⁴⁶ The study revealed, for example, that three judges in this same office awarded benefits to no more than 30 percent of their applicants, while three other judges awarded to more than 70 percent.⁴⁷ Other studies have suggested that racial disparities may exist in Social Security disability awards, with certain Black applicants tending to receive less favorable outcomes than White applicants.⁴⁸ Against this kind of track record, automated systems promise distinct advantages when they can be shown to deliver fairer, more consistent, and even speedier decisions.

But humans will still be good at listening and empathizing with the predicaments of those who are seeking assistance or other decisions from government, or who otherwise find themselves subjected to its constraints.⁴⁹ It is that human quality of empathy that should lead the administrative law of procedural due process to move beyond just its current emphasis on reducing errors and lowering costs.

To some judges, the need for an administrative law of empathy may lead them to ask whether members of the public have a “right to a human decision” within

an automated state.⁵⁰ But not all human decisions are necessarily empathic ones. Moreover, a right to a human decision would bring with it the possibility that the law would accept all the flaws in human decision-making simply to retain one of the virtues of human engagement. If automated decisions turn out increasingly to be more accurate and less biased than human ones, a right to a decision by humans would seem to deny the public the desirable improvements in governmental performance that algorithms can deliver.

Administrative law need not stand in the way of these improvements. It can accept the use of machine learning algorithms while nevertheless pushing government forward toward additional opportunities for listening and compassionate responses.⁵¹ Much as the Supreme Court in *Goldberg v. Kelly* insisted on a pre-termination hearing for welfare recipients, courts in the future can ask whether certain interests are of a sufficient quality and importance to demand that agencies provide supplemental engagement and assistance with individuals subjected to automated processes. Courts could in this way seek to reinforce best practices in agency efforts to provide empathic outreach and assistance.

In the end, if administrative law in an automated state is to adopt any new rights, society might be better served if courts avoid the recognition of a right to a human decision. Instead, courts could consider and seek to define a right to human *empathy*.

AUTHOR'S NOTE

Thank you to Lavi Ben Dor for research assistance and to Richard Berk, Amanda Greene, and Mark Tushnet for helpful comments on an earlier draft.

ABOUT THE AUTHOR

Cary Coglianese is the Edward B. Shils Professor of Law, Professor of Political Science, and Director of the Penn Program on Regulation at the University of Pennsylvania. He is the editor of *Achieving Regulatory Excellence* (2016), *Does Regulation Kill Jobs?* (with Adam M. Finkel and Chris Carrigan, 2013), and *Regulatory Breakdown? The Crisis of Confidence in U.S. Regulation* (2012).

ENDNOTES

¹ See Cary Coglianese and David Lehr, "Regulating by Robot: Administrative Decision Making in the Machine-Learning Era," *The Georgetown Law Journal* 105 (5) (2017): 1147, 1171.

- ² See Cary Coglianese and David Lehr, “Transparency and Algorithmic Governance,” *Administrative Law Review* 71 (1) (2019): 1, 9.
- ³ See Benjamin H. Barton and Stephanos Bibas, *Rebooting Justice: More Technology, Fewer Lawyers, and the Future of Law* (New York: Encounter Books, 2017), 111–115.
- ⁴ See Nathan Heller, “Estonia, the Digital Republic,” *The New Yorker*, December 11, 2017, <https://www.newyorker.com/magazine/2017/12/18/estonia-the-digital-republic>; and Republic of Estonia, Ministry of Economic Affairs and Communications, *Report of Estonia’s AI Taskforce* (Tallinn: Republic of Estonia, Ministry of Economic Affairs and Communications, 2019), https://f98cc689-5814-47ec-86b3-db505a7c3978.filesusr.com/ugd/7df26f_486454c9f32340b28206e140350159cf.pdf.
- ⁵ See Eric Niler, “Can AI Be a Fair Judge? Estonia Thinks So,” *Wired*, March 25, 2019, <https://www.wired.com/story/can-ai-be-fair-judge-court-estonia-thinks-so/>.
- ⁶ See United Nations Department of Economic and Social Affairs, *E-Government Survey 2020: Digital Government in the Decade of Action for Sustainable Development* (New York: United Nations, 2020), [https://publicadministration.un.org/egovkb/Portals/egovkb/Documents/un/2020-Survey/2020%20UN%20E-Government%20Survey%20\(Full%20Report\).pdf](https://publicadministration.un.org/egovkb/Portals/egovkb/Documents/un/2020-Survey/2020%20UN%20E-Government%20Survey%20(Full%20Report).pdf).
- ⁷ See Cary Coglianese and Lavi M. Ben Dor, “AI in Adjudication and Administration,” *Brooklyn Law Review* 86 (forthcoming 2021), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3501067.
- ⁸ See David Freeman Engstrom, Daniel E. Ho, Catherine M. Sharkey, and Mariano-Florentino Cuéllar, *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies* (Washington, D.C.: The Administrative Conference of the United States, 2020), <https://www-cdn.law.stanford.edu/wp-content/uploads/2020/02/ACUS-AI-Report.pdf>.
- ⁹ See Gabe Cherry, “Google, U-M to Build Digital Tools for Flint Water Crisis,” *University of Michigan News*, May 3, 2016, <https://news.umich.edu/google-u-m-to-build-digital-tools-for-flint-water-crisis/>; City of Chicago, “Food Inspection Forecasting: Optimizing Inspections with Analytics,” <https://chicago.github.io/food-inspections-evaluation/> (accessed November 21, 2020); Robert Sullivan, “Innovations in Identifying People Who Frequently Use Criminal Justice and Healthcare Systems,” *Policy Research Associates*, May 16, 2018, <https://www.prainc.com/innovations-identification-cj-healthcare/>; Harvard Kennedy School, Data-Smart City Solutions, “A Catalogue of Civic Data Use Cases: How Can Data and Analytics Be Used to Enhance City Operations?” October 9, 2019, <https://datasmart.ash.harvard.edu/news/article/how-can-data-and-analytics-be-used-to-enhance-city-operations-723>; and University of Pennsylvania, “Uses in Government,” <https://www.law.upenn.edu/institutes/ppr/optimizing-government-project/government.php#municipal> (accessed November 21, 2020).
- ¹⁰ Dan Hurley, “Can an Algorithm Tell When Kids Are in Danger?” *The New York Times Magazine*, January 2, 2018, <https://www.nytimes.com/2018/01/02/magazine/can-an-algorithm-tell-when-kids-are-in-danger.html>.
- ¹¹ See Coglianese and Ben Dor, “AI in Adjudication and Administration.”
- ¹² See Tim Lau, “Predicting Policing Explained,” *Brennan Center for Justice*, April 1, 2020, <https://www.brennancenter.org/our-work/research-reports/predictive-policing-explained>.

- ¹³ See Shirin Ghaffary and Rani Molla, “Here’s Where the U.S. Government Is Using Facial Recognition Technology to Surveil Americans,” Vox, December 10, 2019, <https://www.vox.com/recode/2019/7/18/20698307/facial-recognition-technology-us-government-fight-for-the-future>.
- ¹⁴ See Tod Newcombe, “Aiming Analytics at Our \$3.5 Billion Unemployment Insurance Problem,” Government Technology, March 2017, <https://www.govtech.com/data/Aiming-Analytics-at-Our-35-Billion-Unemployment-Insurance-Problem.html>.
- ¹⁵ See David Freeman Engstrom and Daniel E. Ho, “Algorithmic Accountability in the Administrative State,” *Yale Journal on Regulation* 37 (3) (2020): 800, 815–819; and Richard Rubin, “AI Comes to the Tax Code,” *The Wall Street Journal*, February 26, 2020, <https://www.wsj.com/articles/ai-comes-to-the-tax-code-11582713000>.
- ¹⁶ Partnership for Public Service and IBM Center for the Business of Government, *More Than Meets AI: Assessing the Impact of Artificial Intelligence on the Work of Government* (Washington, D.C.: Partnership for Public Service and IBM Center for the Business of Government, 2019), <https://ourpublicservice.org/wp-content/uploads/2019/02/More-Than-Meets-AI.pdf>.
- ¹⁷ Such tenets are reflected in both the notion of due process as well as the general standard that agency action should not be arbitrary and capricious.
- ¹⁸ See Coglianese and Lehr, “Transparency and Algorithmic Governance,” 26–29.
- ¹⁹ 462 U.S. 87, 103 (1983).
- ²⁰ 490 U.S. 360, 371 (2011), quoting *Kleppe v. Sierra Club*, 427 U.S. 390, 412 (1976).
- ²¹ See Coglianese and Lehr, “Transparency and Algorithmic Governance,” 49–55.
- ²² 463 U.S. 29 (1983).
- ²³ Validation, which should take place before abandoning the status quo of a human-based process, could involve testing the algorithm on randomly selected cases that are also, in tandem, decided by humans following normal procedures. Closer scrutiny could be provided by panels of human experts of discrepancies between the results of digital systems and the initial human decision-makers.
- ²⁴ This regulatory review regimen is outlined in Executive Order 12866 of September 30, 1993, *Federal Register* 58 (190) (1993).
- ²⁵ 467 U.S. 837 (1984). As a formal matter, the Court grounded *Chevron* deference in an explicit or implicit delegation of clarifying or gap-filling authority to the agency. See Cary Coglianese, “*Chevron*’s Interstitial Steps,” *The George Washington Law Review* 85 (5) (2017): 1339, 1347–1351.
- ²⁶ 467 U.S. 837, 865.
- ²⁷ Benedict Carey, “Can an Algorithm Prevent Suicide?” *The New York Times*, November 23, 2020, <https://www.nytimes.com/2020/11/23/health/artificial-intelligence-veterans-suicide.html>.
- ²⁸ *Ibid.*
- ²⁹ For a discussion of the intelligible principle doctrine, see Cary Coglianese, “Dimensions of Delegation,” *University of Pennsylvania Law Review* 167 (7) (2019): 1849.
- ³⁰ See *Washington v. Davis*, 426 U.S. 229, 239 (1976).

- ³¹ See *Fisher v. University of Texas at Austin*, 133 S. Ct. 2411 (2013); *Grutter v. Bollinger*, 539 U.S. 306 (2003); *Gratz v. Bollinger*, 539 U.S. 244 (2003); and Coglianese and Lehr, “Regulating by Robot.”
- ³² For an accessible account of technical aspects of algorithmic fairness, see Michael Kearns and Aaron Roth, *The Ethical Algorithm: The Science of Socially Aware Algorithm Design* (New York: Oxford University Press, 2019), 57–93. For a cogent discussion of how digital algorithms can generate results more fair than human processes, see Sandra G. Mayson, “Bias In, Bias Out,” *Yale Law Journal* 128 (8) (2019): 2218, 2277–2281.
- ³³ See, for example, Richard A. Berk and Arun Kumar Kuchibhotla, “Improving Fairness in Criminal Justice Algorithmic Risk Assessments Using Conformal Prediction Sets,” (2021), arXiv:2008.11664 [stat.AP]; James E. Johndrow and Kristian Lum, “An Algorithm for Removing Sensitive Information: Application to Race-Independent Recidivism Prediction,” *The Annals of Applied Statistics* 13 (1) (2019): 189; and Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Cass R. Sunstein, “Discrimination in the Age of Algorithms,” *Journal of Legal Analysis* 10 (2018): 113.
- ³⁴ See, for example, Danielle Keats Citron, “Technological Due Process,” *Washington University Law Review* 85 (6) (2008): 1249; and Karen Yeung, “Algorithmic Regulation: A Critical Interrogation,” *Regulation and Government* 12 (4) (2018): 505.
- ³⁵ For example, norms will surely develop about how agencies should document their choices in designing algorithmic systems. See, for example, Timnit Gebru, Jamie Morgenstern, Briana Vecchione, et al., “Datasheets for Datasets” (2020), <https://arxiv.org/abs/1803.09010>.
- ³⁶ Channeling Voltaire, if *Baltimore Gas & Electric* did not exist, courts might still find it necessary to invoke its deference. See Adrian Vermeule, *Law’s Abnegation: From Law’s Empire to the Administrative State* (Cambridge, Mass.: Harvard University Press, 2016).
- ³⁷ As Steven Appel and I have noted elsewhere, “it is not hard to imagine a time in the near future when the public actually comes to expect their public servants to rely on such technologies. As complex machine-learning algorithms proliferate in the private sector, members of the public may well come to expect similar accuracy and automated services from their governments.” Steven M. Appel and Cary Coglianese, “Algorithmic Governance and Administrative Law,” in *Cambridge Handbook on the Law of Algorithms: Human Rights, Intellectual Property, Government Regulation*, ed. Woodrow Barfield (Cambridge: Cambridge University Press, 2021), 162, 165.
- ³⁸ Notably, technologists are even exploring the possibility of building empathy into automated tools. See, for example, Pascale Fung, Dario Bertero, Yan Wan, et al., “Towards Empathetic Human-Robot Interactions,” in *Computational Linguistics and Intelligent Text Processing*, ed. Alexander Gelbukh (New York: Springer International Publishing, 2018).
- ³⁹ Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order* (New York: Houghton Mifflin Harcourt, 2018), 210.
- ⁴⁰ The advent of an automated administrative state will unquestionably lead to changes in the government labor force, much as the expanded use of artificial intelligence in the private sector will lead to changes in the labor market more generally. David Autor and Anna M. Salomons, “Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share,” *Brookings Papers on Economic Activity*, March 8, 2018; and David Autor, David Mindell, and Elisabeth Reynolds, *The Work of the Future: Building Better Jobs in an Age of Intelligent Machines* (Cambridge, Mass.: MIT Work of the Future, 2020),

<https://workofthefuture.mit.edu/wp-content/uploads/2021/01/2020-Final-Report4.pdf>. The shift to a government workforce increasingly organized around empathy harkens to a broader shift from a manufacturing economy to service economy, or perhaps still further to a “sharing” economy. See Cary Coglianese, “Optimizing Regulation for an Optimizing Economy,” *University of Pennsylvania Journal of Law and Public Affairs* 4 (1) (2018): 3–4.

- ⁴¹ Partnership for Public Service and IBM Center for the Business of Government, *More Than Meets AI*, 8.
- ⁴² Allen E. Lind and Tom Tyler, *The Social Psychology of Procedural Justice* (New York: Springer International Publishing, 1988).
- ⁴³ Amanda Greene, “Competence, Fairness, and Caring: The Three Keys to Government Legitimacy,” Centre for Public Impact, February 27, 2018, <https://www.centreforpublicimpact.org/the-three-keys-government-legitimacy/>.
- ⁴⁴ 397 U.S. 254 (1970). See Lucie E. White, “Subordination, Rhetorical Survival Skills, and Sunday Shoes: Notes on the Hearing of Mrs. G.,” *Buffalo Law Review* 38 (1) (1990).
- ⁴⁵ 424 U.S. 319 (1976).
- ⁴⁶ “Social Security Awards Depend More on Judge Than Facts,” TRAC Social Security Administration, July 4, 2011, <https://trac.syr.edu/tracreports/ssa/254/>. The Social Security Administration sharply disputed aspects of this study.
- ⁴⁷ *Ibid.*
- ⁴⁸ See, for example, U.S. General Accounting Office, “Racial Difference in Disability Decisions Warrants Further Investigation,” B-247327 (Washington, D.C.: U.S. General Accounting Office, 1992), <https://www.gao.gov/assets/160/151781.pdf>; and Erin M. Godtland, Michele Grgich, Carol Dawn Petersen, et al., “Racial Disparities in Federal Disability Benefits,” *Contemporary Economic Policy* 25 (1) (2007): 27.
- ⁴⁹ For a discussion of the importance of empathy in the exercise of regulatory authority, see Cary Coglianese, *Listening, Learning, and Leading: A Framework for Regulatory Excellence* (Philadelphia: University of Pennsylvania Law School, 2015), 23–25, <https://www.law.upenn.edu/live/files/4946-pprfinalconvenersreport.pdf>.
- ⁵⁰ For an excellent treatment of this question, see Aziz Z. Huq, “A Right to a Human Decision,” *Virginia Law Review* 106 (3) (2020): 611.
- ⁵¹ Sometimes the compassionate response may even call for overriding an automated decision: that is, to have a human official exhibit mercy and reach a different decision on an individual basis. After all, automated systems themselves will still result in errors, and joint human-machine systems may well at times do better to reduce errors than either humans or machines operating separately. The challenge, though, will be to ensure enough structure around the discretion to override automated outcomes, lest human exceptions come to swallow automated rules. See Cary Coglianese, Gabriel Schefler, and Daniel E. Walters, “Unrules,” *Stanford Law Review* 73 (4) (2021): 885. One solution might be to create automated systems specifically designed to help with this very problem. If an automated system generates not only an outcome but also an estimate of confidence in that outcome, humans may be guided to go beyond empathic listening and deliver merciful exceptions only in those instances where a system’s estimated confidence is sufficiently low.