



ARTIFICIAL INTELLIGENCE AND THE ADMINISTRATIVE STATE

Mariano-Florentino Cuéllar[†]

We rely on agencies to increase air quality and mitigate climate change, protect public health and safety, and safeguard the integrity of financial markets. Nearly a century ago Max Weber cogently observed that the modern nation-state depends on bureaucracy—or, in modern parlance, on administrative agencies.¹ He would not have been surprised to see how, even as self-driving cars navigate the streets of Pittsburgh, Pennsylvania and Mountain View, California, agencies staffed by bureaucrats and overseen by administrators have remained the essential organizational technology of the administrative state.² Whether those agency administrators exercise sufficient independent judgment as individuals to warrant the integrity and accountability of a decision has, in turn, been the subject of some classic administrative law cases.³

But technological change is creating new dilemmas and opportunities for the administrative state. Agencies today can rely on sophisticated computer programs—programs that agencies could use to make or support decisions, and that could therefore assume an increasingly prominent role in the regulatory process. The smartphones that so many Americans carry around in their pockets are far more powerful—and an order of magnitude cheaper—than the vast computers scientists and the military used a generation ago. In the coming years, computing power and storage will almost certainly become even cheaper, surveillance more pervasive, software architecture more flexible, and the limitations of human decision-makers will become more salient.

[†] Justice on the California Supreme Court and Visiting Professor at Stanford Law School and Harvard Law School. This essay is an edited version of the author's 2016 Distinguished Lecture on Regulation at the University of Pennsylvania Law School.

¹ MAX WEBER, *ECONOMY AND SOCIETY* (1922), https://archive.org/stream/MaxWeberEconomyAndSociety/MaxWeberEconomyAndSociety_djvu.txt.

² Liz Reid, *What It's Like to Ride in a (Nearly) Self-Driving Uber*, NPR (Sept. 14, 2016), <https://www.npr.org/sections/alltechconsidered/2016/09/14/493823483/self-driving-cars-take-to-the-streets-of-pittsburgh-courtesy-of-uber>.

³ *Heckler v. Campbell*, 461 U.S. 458 (1983); *Morgan v. United States*, 304 U.S. 1 (1938).

Traditional expert systems used law-like techniques to search through potential options when analyzing how to diagnose certain medical conditions, or how to categorize a particular kind of molecule—but they were cumbersome at best when it came to some of the seemingly simplest things that people could do almost “without thinking,” like classifying visual objects, interpreting idiomatic expressions, or decoding nonverbal communication. As computing power gets cheaper and software improves, expert systems are ever more able to sift through millions of options quite quickly. But an even bigger change is underway in the realm of so-called “machine learning,” where the software architecture uses two interesting techniques.

One technique involves so-called “neural networks,” which are inspired by the layout of the human brain to spot patterns and leverage “big data.” “Deep learning” systems embody a particular architecture for neural networks that avoid some persistent problems neural networks have had in developing adaptive responses to new data and have sparked particular interest because of its capacity to solve pattern-recognition problems in computer vision and other fields.

Meanwhile, “genetic algorithms” that emerge by developing simple algorithms—or baby computer programs—to solve a problem like spotting suspicious financial transactions, allowing those algorithms to mutate slightly over time, and then selecting for the algorithms that beat the others on a given metric. It is a great way to write a nearly security-proof, pesky computer virus, which we will return to later in this series. But it is more generally through machine learning that new progress is underway on many of those apparently simple but devilishly hard technical problems, like vision and speech recognition.

Because of these changes, lawyers for regulated industries, citizens facing possible search or arrest from the police, and individuals seeking asylum will find themselves interacting with agency officials who heavily rely on software to make decisions—or perhaps these members of the public will be interacting directly with the software itself. More extensive use of programs designed to supplement—or even replace—human decision-makers will become commonplace as computing power and memory become cheaper, data from surveillance become more pervasive, and economic and military pressures drive adoption. The public already relies on software to recommend romantic partners and investments. As autonomous and elaborate decision-support programs become more common, social norms will continue to change about the propriety of relying on computers to make decisions. Although computer programs analyzing vast amounts of information may hold some promise for making better use of data, enhancing transparency, and reducing inconsistency in bureaucratic justice, such reliance may bring about subtle consequences and deeper questions that merit careful scrutiny.

What should we make of a world where the entities entrusted to exercise administrative power are not agencies but software programs that leverage

the fast-developing technology of artificial intelligence? Imagine a series of sleek black boxes—capable of sustaining a cogent conversation with an expert, and networked to an elaborate structure of machines, data, and coded instruction sets—that deliver bureaucratic justice. It could begin innocently enough, with anodyne decision-support programs for administrative law judges adjudicating disability claims, or for hearing examiners at the U.S. Environmental Protection Agency. But as the interfaces became more intuitive and the analytical capacity more sophisticated, the black boxes might steadily climb up the bureaucratic ladder, displacing supervisors, division heads, and even agency administrators. All of which could recast—or even disrupt—legally-sanctioned bureaucratic authority.

It may seem simple enough to determine the expected value of these changes in social welfare terms. Consider the choice, for example, to replace an administrative law judge working on disability determinations, or even an Under Secretary responsible for food safety, with an expert system—one that could replay in exquisite detail the sequence of decision rules it relied on to render a judgment. Any reasonable effort to judge the quality of that judgment mainly depends on how a statute or regulatory rule defines a domain-specific metric of success. Because such delegation could affect variables that cut across domains—such as perceptions of government legitimacy, cybersecurity risks, and the extent of presidential power—even more important would be an uncontroversial metric of social welfare, along with certain assumptions to minimize the difficult trade-offs across domains.

But more profound challenges would arise in the myriad situations where the unambiguous metric is not so easily available. Think about the subtle choices involving drug approval, asylum applications, bioethics, and the protection of endangered species. In all of these areas, heavy reliance on artificially intelligent systems could also make it harder for lawmakers, courts, and the public to assess the consequences of automated agency decision-making where the trade-offs are complex.

We may ultimately find that the choices we make about automation will be part of a broader conflict about the role of people in an economy that sheds a large proportion of existing job categories more quickly than expected, even as it continues to enhance automation technologies that humans find, like a sweet-tasting artificial strawberry dessert, occasionally more satisfying than the “natural” alternative. As these questions become more familiar, the administrative state will continue confronting a host of challenges entirely recognizable to Weber—from striking the right balance between agency insulation and responsiveness to the role of tradition in bureaucratic decision-making. But increasingly, the dilemma agencies and the public will face is what to do about the aforementioned sleek black boxes that promise to make governing far simpler and cheaper. Whether those boxes also give us an accurate account of who gains or loses in the process is not something we should take for granted.

I. THE SURPRISING USE OF AUTOMATION BY REGULATORY AGENCIES

Let us start by acknowledging that humans make mistakes. Social psychologists, economists, political scientists, and even policymakers routinely acknowledge the limitations of how humans tend to consider probabilities, or otherwise weigh the consequences of particular decisions.⁴ Decision-makers may exhibit racial or gender biases, may over- or underweigh the importance of a particular piece of information, may naively assume their own wisdom, or may insist on the naiveté of rivals. Even thoughtful experts who are familiar with the subtleties of environmental, national security, or public health data may fail to recognize patterns that can give agencies useful ideas about how to achieve their responsibilities.

It is certainly understandable, then, why societies could become interested in making greater use of computer systems that hold the promise of improving the quality and integrity of administrative decisions. Government agencies are beginning to rely more on computer programs to make decisions, and this trend will likely accelerate. An example involving federal regulation of pesticides highlights the subtle ways in which computer-based analysis and legal standards could interact—as well as the reasons why agencies may embrace new analytical techniques that heavily rely on automation.

The U.S. Environmental Protection Agency (EPA) administers the Federal Insecticide, Fungicide, and Rodenticide Act, which requires the registration of pesticides before marketing in interstate or foreign commerce.⁵ The current toxicity testing for pesticides depends heavily on assessing animals' reactions to chemicals—a technique that can be easily criticized as costly, slow, and inhumane. At the most basic level, current toxicity testing methods limit the number of chemicals the EPA can test, even though it faces strong pressures to test more than 80,000 chemicals. But further, it limits the number of toxicity pathways one can test, the levels of biological organizations one can examine, the range of exposure conditions one can consider, and the life stages, genders, and species one can cover.

Given the inadequacy of current methods of toxicity research, the National Academy of Sciences published a report in 2007 calling for a transformative shift in toxicity and risk assessment and increased use of computational toxicology.⁶ In response, the EPA is introducing many different forms of computational methods in regulating pesticides.

⁴ See MARK KELMAN, *THE HEURISTICS DEBATE* (1st ed. 2011).

⁵ 7 U.S.C. §136 et seq. (1996). See generally NAT'L ACAD. SCIS., *TOXICITY TESTING IN THE 21ST CENTURY: A VISION AND A STRATEGY* (2007).

⁶ See NAT'L ACAD. SCIS., *supra* note 5; Brinda Mahadevan et al., *Genetic Toxicology in the 21st Century: Reflections and Future Directions*, NAT'L CTR FOR BIOTECHNOLOGY INFO. (Apr. 28, 2011), <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3160238/>; Robert J. Kavlock et al., *Computational Toxicology—A State of the Science Mini Review*, 103 TOXICOLOGICAL SCIS. 14 (2008).

Computation is also helping the EPA better calculate and predict environmental exposure to chemicals. Modern computational methods can build complex models that consider many variables that determine the level of exposure to toxic chemicals, such as the difference in exposure to animals versus humans, variability in exposure to humans, and the overall uncertainty of these predictions.⁷

To support these efforts, the EPA is involved in massive data collection. It created the Aggregated Computational Toxicology Resource, a relational data warehouse for chemical and toxicity data from various public sources to support data mining and modeling.⁸ The EPA is also poised to start using virtual tissues; the agency is currently developing a “virtual liver” at the EPA’s National Center for Computational Toxicology.⁹

The EPA’s reliance on computational toxicology underscores how agency decisions may increasingly implicate not only human choices about research methods, but architectural choices in the development of algorithms and neural networks to analyze data in new ways.

Changes in disability claims, too, may emerge as agencies seek to resolve logistical problems while compensating for inconsistencies of human judgment. In 2013, in an effort to reduce its reliance on paper records, to increase consistency across cases, and to automate some of its workflow, the U.S. Department of Veterans Affairs launched a computerized case management system for incoming disability claims.¹⁰ The software reportedly automates how the Department determines the level of different veterans’ disabilities for purposes of compensation.¹¹ And importantly, it “calculates the level of disability—from zero to 100%—solely on the vet’s symptoms from the [self-reporting] questionnaire.”¹² In essence, the software took over this responsibility for determining levels of disability from Department “raters”—human beings charged with determining a claimant’s entitlements.

Consider one additional example of automation, from a domain of responsibility shared by the public and private sectors: the testing of pharmaceutical products. As part of its review of new drug applications, the U.S. Food and Drug Administration (FDA) often considers “Population Pharmacokinetics” models, which test how drugs will interact with different bodies, depending on age, weight, and other factors. Traditionally, experts

⁷ See HERBERT A. SIMON, *ADMINISTRATIVE BEHAVIOR* (4th ed. 1997).

⁸ See Richard Judson, et al., *ACToR—Aggregated Computational Toxicology Resource*, 233 *TOXICOLOGY & PHARMACOLOGY* 7 (2008).

⁹ *About the National Center for Computational Toxicology (NCCT)*, ENV’T. PROT. AGENCY, https://19january2017snapshot.epa.gov/aboutepa/about-national-center-computational-toxicology-ncct_.html.

¹⁰ Bob Brewin, *Goodbye Paper: VA Installs Automated Claims System in All Regional Offices*, NEXTGOV (June 17, 2013), <https://www.nextgov.com/cio-briefing/2013/06/goodbye-paper-va-installs-automated-claims-system-all-regional-offices/65030/>.

¹¹ Daniel Huang, *Automated System Often Unjustly Boosts Veterans’ Disability Benefits*, WALL ST. J., May 11, 2015, <https://www.wsj.com/articles/automated-system-often-unjustly-boosts-veterans-disability-benefits-1431387826>.

¹² *Id.*

known as “pharmacometricians” would select several hundred statistical models (not real people) on which to test these drug interactions. As expected, choosing which models to include was time consuming and labor intensive.¹³

As an alternative, the FDA recently approved a new drug application in which models were selected by an algorithm. According to the developer’s press release announcing the fact, such “automated model selection provides pharmaceutical and biotech companies results in less than half the time and at a lower cost compared to the traditional method.”¹⁴

As these several examples suggest, greater reliance on artificial intelligence has much that will appeal to government officials. In the years ahead, government contractors will push new technologies to sell to administrative agencies. Outside lawyers will continue to criticize arbitrary agency decisions. Civil society groups will make the case for more predictable and analytically-sound administrative decisions. Taken together, these various pressures are likely to encourage agencies to find ways of relying on data and computer programs to make regulatory decisions.

And the promise of automation in the administrative state will not be entirely contingent on computer systems that mimic human interaction. Some travelers may prefer to be screened by even a fairly conventional computer system, rather than by an agent whose biases and limitations could color her judgment. After all, human decision-makers get things wrong.

The use of statistical and other predictive techniques by computers could improve not only individual decisions, but systemic bureaucratic performance as well. As computing technology improves, new possibilities will emerge to juxtapose two seemingly opposite qualities that could make automation more difficult to resist—the ability to analyze data and make predictions in subtle fashion that does not easily track human intuition, coupled with the capacity to make increasingly persuasive arguments to defend a decision.

But what, exactly, could more robust reliance on sophisticated information technologies accomplish? The simplest scenario is one where information technology duplicates what a human administrator could do, at a lower cost. Alternatively, the right expert systems could also screen out biases and certain heuristics that are considered, in the aggregate, to be undesirable, such as availability and vividness heuristics.

Even more intriguingly, computer programs could make it possible for government officials to analyze information for the purpose of predicting outcomes or responding to potential strategic behavior in a fashion that would be enormously difficult—if not impossible—for a human decision-maker to approximate. Massive concentrations of data analyzed by neural networks could generate intricate new predictions of how criminal enterprises, for example, adjust to new anti-money laundering measures, and

¹³ *Automated Model Selection*, NUVENTRA, <http://www.nuventra.com/services/darwin/>.

¹⁴ *FDA Approves First NDA Using Nuventra’s Software for Automated Population PK Model Selection – DARWIN*, NUVENTRA (Nov. 2, 2014), <http://www.nuventra.com/fda-darwin-population-pk-automated-model-selection/>.

what mix of counter-measures could help neutralize new forms of subterfuge to hide money corruptly stolen from foreign governments or obtained through fraud. Machine learning techniques could help food safety administrators further target scarce inspection resources to conduct the limited number of foreign inspections that are possible given existing resource constraints. These possibilities make it hard to ignore the opportunities for automating certain aspects of the administrative state—and all the more important to consider the normative questions that the uses of automation will raise.

II. PREPARING FOR CYBERDELEGATION AND ITS RISKS

The U.S. Environmental Protection Agency (EPA), the U.S. Department of Veterans Affairs, and the U.S. Food and Drug Administration are just a few of the agencies turning to automation as a way to improve regulatory functioning. In the years ahead, we will see only more instances of agency use of cyberdelegation—or the reliance on computer programs to make government decisions. Thoughtful use of computers in administrative government—and in particular deployment of artificial intelligence technologies involving expert systems and deep learning—have the potential to increase consistency in decision-making and to help agency officials understand a complex and changing world to make better decisions.

But the advantages of cyberdelegation in the administrative state will bring with them at least four sets of challenges warranting careful scrutiny. First, the societal value of government reliance on computer programs will depend on highly contestable assessments of programs' objectives. And deciding how to instruct computer programs on matters of broad public concern—and telling them what to maximize—will be more difficult in practice than in theory.

These difficulties will arise even when there is widespread societal agreement about a given general goal, such as keeping food safe at a reasonable cost, or reducing vulnerability to terrorist attacks, in part because agreement at a high level of generality rarely translates into consensus on how to implement policies through administrative agencies.

Plenty of debate will occur within agencies and among legislators about the precise mix of goals that should animate various administrative decisions, such as the imposition of economic sanctions. It is easy enough to suggest that the goal is to change the behavior of the target country. But the details matter. Often, implementing policy involves political tradeoffs that an expert system could elide but would still be making, implicitly, by applying a particular analytical technique.

A second challenge will be determining how much will be lost when human cognition is replaced by machines. Our often under-theorized goals must inform whether we should try to screen out features of human cognition—including the often-mentioned “heuristics” and “biases”—that diverge from conventional and easily systematized accounts of rationality.

There is no reason to think that all heuristics and biases are bad from a social welfare perspective. Whether a heuristic is valuable depends on what goal one has for society.

Some features of human cognition that vary from conventional rationality—such as the tendency to weigh more heavily the stories of specific individuals rather than aggregate statistical information—may be integral to qualities such as empathy, or to the ability of policymakers to explain governmental decisions to the public.¹⁵ Accordingly, at least in some circumstances, quirks of human decision-making that are often treated as “biases” to be screened out by computer algorithms may instead merit an increasingly important place in legal decision-making as many routine decisions are guided by algorithms.

Third, potential side effects from automation must be considered. The incorporation of computer programs into the administrative state could carry with them cybersecurity risks and have other adverse impacts that will not necessarily be weighed in a calculus that may encourage reliance on computer programs.

It may be tempting to ignore cybersecurity problems because we have yet to develop an effective technical means for quantifying the risks. But it would be a serious mistake to consider the benefits of automation without considering the associated security problems.

For example, greater EPA reliance on pervasive data gathering and computer programs to target enforcement could result in a world with less pollution, but also one more vulnerable to cybersecurity threats that could, at a minimum, undermine the integrity of the regulatory process and, at worst, exploit vulnerabilities to undermine industrial infrastructure. Cybersecurity problems should loom especially large due to the many examples of governmental failures involving information technology.¹⁶

Fourth, heavy reliance on computer programs may adversely affect the extent of deliberation that occurs in the administrative state. Implicit in democratic governance is an aspiration for dialogue and exchange of reasons that are capable of being understood, accepted, or rejected by policymakers, representatives of organized interests, and members of the public.

Except when computerized decisions can rely on relatively straightforward, rule-like structures, difficulties will arise in supplying explanations of how decisions were made that could be sufficiently understood by policymakers and the public.¹⁷ For example, if computer systems determined how to allocate scarce inspection resources among processing facilities handling the increasing proportion of the American food

¹⁵ See KELMAN, *supra* note 4; CASS R. SUNSTEIN, *SIMPLER: THE FUTURE OF GOVERNMENT* (2013).

¹⁶ Niam Yaraghi, *Doomed: Challenges and Solutions to Government IT Projects*, BROOKINGS INST. (Aug. 25, 2015), <https://www.brookings.edu/blog/techtank/2015/08/25/doomed-challenges-and-solutions-to-government-it-projects/>.

¹⁷ John Markoff, *Google Car Exposes Regulatory Divide on Computers as Drivers*, N.Y. TIMES, Feb. 10, 2016, https://www.nytimes.com/2016/02/11/technology/nhtsa-blurs-the-line-between-human-and-computer-drivers.html?_r=0.

supply that comes from abroad, it would probably matter to importers and consumers that these systems would be unable to yield carefully-reasoned explanations for the choices undertaken.

Confronting these four major challenges today is important because major path-dependent effects will make it difficult to undo the use of algorithms once they are incorporated into legal decision-making. Path dependence will arise because infrastructure is costly to replace and habituates people to make decisions in a particular way.

For example, given recent advances in DNA sequencing and genetic medicine, it is not difficult to envision an ever-greater role for expert systems in analyzing information relevant to the approval of specialized drugs.¹⁸ Even though computer programs and organizational expertise may function as complements today, they may become substitutes at a later time. Once an agency's organizational expertise begins to erode due to greater reliance on computerized decision systems, the agency will face steep costs in recovering that expertise.

Overall, the administrative state is about expertise and, more importantly, its translation and engagement with the broader public. Administrative decision-making involves moving back and forth from discourses surrounding expert knowledge and legal authority to conversations that entail public deliberation and moral debate. The core question underlying cyberdelegation will be what happens to this process of translation when automated systems have a more prominent role in the administrative state.

This is not to say that the status quo is any deliberative panacea. On the contrary, it is easy to criticize the current administrative state for its lack of opportunities to allow the public to participate in decisions. Yet the growing reliance on automated computer programs to make sensitive decisions in the administrative state will only complicate what little deliberation *does* occur. Cyberdelegation risks diffusing responsibility between the agency's leadership and the team or set of machines that designed the relevant software, raising the likelihood that decisions would be made on a basis that is different from what could be understood or even explained by human participants.

III. DECIDING WHETHER SOFTWARE WILL EAT THE BUREAUCRACY

With widely circulating media accounts that a foreign power used cyber-intrusions in an effort to affect a recent American national election, it is not radical to suggest that reliance on computers to make agency decisions is a risky enterprise.¹⁹ But in some ways, cybersecurity problems are just the tip of the iceberg. From cybersecurity risks to changes in public

¹⁸ *Precision Medicine Initiative*, U.S. FOOD & DRUG ADMIN. (Oct. 3, 2016), <https://web.archive.org/web/20161219122430/http://www.fda.gov/ScienceResearch/SpecialTopics/PrecisionMedicine/default.htm>.

¹⁹ Mark Mazzetti & Eric Lichtblau, *C.I.A. Judgment on Russia Built on Swell of Evidence*, N.Y. TIMES, Dec. 11, 2016, <http://www.nytimes.com/2016/12/11/us/politics/cia-judgment-intelligence-russia-hacking-evidence.html>.

deliberation, government agencies' use of automation and artificial intelligence will pose numerous challenges for the administrative state.

Although no simple compass or rubric exists to decide precisely how to navigate these uncharted waters, the following ideas offer a few possibilities for how agencies, policymakers, and the courts could help increase society's capacity to make informed choices about the use of automation in the administrative state.

First, it may be worth exploring how we may better police the extent of human decision-maker engagement with automated expert systems. Until now, the courts have been reluctant to probe the actual decision-making of administrative leaders under the so-called presumption of regularity that emerged over time following *Morgan v. United States*.²⁰ In rejecting a challenge to an order by the Secretary of Agriculture fixing maximum rates to be charged by market agencies at the Kansas City Stockyards, the Supreme Court in *Morgan* declined to allow an intrusive analysis of the Secretary's actual decision-making process and considerations. "It was not the function of the court to probe the mental processes of the secretary in reaching his conclusions," the Court concluded, "if he gave the hearing which the law required."²¹ With courts loath to stray from this presumption of regularity over the decades, it has persisted—and with it, courts' unwillingness to police exactly *by whom* a decision is taken.²²

As reliance on information technology increases, courts and policymakers should consider taking more seriously requiring accountability to be lodged in specific decision-makers. Perhaps it is time to consider recalibrating the "presumption of regularity"—to ensure that agency officials have clearly recognized the risks of relying on automated analytical techniques that are too complex or opaque for officials themselves to understand entirely.

As a practical matter, this approach raises difficult further questions about the scope of discovery in suits to review administrative action, but perhaps those questions are worth facing, given the risk that decision-makers will rely on algorithms they do not fully understand.

Second, on a related note, arbitrary and capricious review may prove most meaningful if it encompasses whether there is consistency between substantive explanations offered in, say, justifications for rulemaking and the analytical techniques actually used to make decisions. It is one thing to justify a program to freeze assets associated with organizations that meet a specific, statutorily grounded threshold of suspicion; it is quite another to deploy algorithms that entirely redefine that threshold, dynamically, in response to new information. Attention to cybersecurity risks may also fit within the context of arbitrary and capricious review.

Third, agencies must accelerate efforts to engage scholars, civil society, and other stakeholders in increasing our understanding of how to harness the

²⁰ *Morgan*, 304 U.S. at 1.

²¹ *Id.*

²² *Postal Service v. Gregory*, 534 U.S. 1 (2001).

analytical capacity of automated computer systems without eroding our sense of how decisions are made. As part of this process, agencies should consider engaging in medium-to-long-term planning about how they would address the use of automation within the rulemaking process. The U.S. Food and Drug Administration could further investigate how trends in artificial intelligence could change the agency's use of outside experts in the drug approval process. Officials at the U.S. Department of Labor may face unexpected challenges arising from labor market changes driven by automation. Virtually all agencies will benefit from explicitly experimenting with different models of decision-making that aim to leverage artificial intelligence technologies while keeping humans in the loop.

These efforts will matter because, increasingly, agencies and entire governments will face the challenge of how to instruct complex machines that will work across domains and agency jurisdiction, aggregate data, and guide human decisions. Government agencies seem to face trouble even when updating conventional information technology infrastructure, so the ability to integrate artificial intelligence into administrative tasks may seem far-fetched.²³

Yet ironically, such weakness could strengthen the case for using systems that adapt and learn. Such systems may prove crucial to reducing the gap between a machine's capacity and that of a person familiar with an agency's culture and organizational routines. As a general matter, as computer systems that perform administrative tasks become adaptive and capable of modifying themselves, the more they are likely to avoid the problems of efficacy and cost that sometimes plague government information technology projects.

But as software becomes more analytically sophisticated, and in particular, more adaptive to the point of being able to rewrite much of its own code, it will be more difficult to predict longer-term consequences ranging from subtle changes in function to unexpected rapid growth of analytical capacity. As machines become more capable of optimizing to achieve the goals we articulate, higher stakes attach to how we articulate those goals and the trade-offs we allow. Crucial to our ability to navigate these dilemmas will be a cadre of lawyers and policymakers who understand artificial intelligence, its possibilities and limits, and particularly its capacity to adapt in unexpected ways.

Lawyers and policymakers will almost certainly need to adjust their approaches to using automation in the administrative state, since different scenarios involving automation are possible, and some will prove far more difficult to manage than others. What makes little sense is to ignore the dilemmas that society will confront as the administrative state comes increasingly to rely on automated systems. Nor is it justified to assume that human decision-making is so fundamentally flawed that it must be tamed by computer system.

²³ U.S. GOV'T ACCOUNTABILITY OFF., GAO-16-696T, INFORMATION TECHNOLOGY: FEDERAL AGENCIES NEED TO ADDRESS AGING LEGACY SYSTEMS (2016) (testimony of David A. Powner, Director, Information Technology Management Issues).

At its core, the administrative state is about reconciling calculations of social welfare with procedural constraints. It is an enterprise that pivots in subtle and profound ways on human institutions, assumptions, and aspirations—however imperfectly fulfilled—for deliberation.

An alternative that promises to make the regulatory process eminently more tractable, technically precise, and less messy by leaning on algorithms and neural networks will likely remain alluring because collective human decisions are as messy and imperfect as human societies are themselves. The biggest risk associated with automation is to assume that most of what concerns the administrative state can be made simpler, more predictable, cheaper, and more effective without any trade-offs. Whether that perspective originates from a deep-seated view that governing is simple or from the seemingly anodyne choices made by a software engineer deciding how to visually present the results of a complex deep learning algorithm, the problem with that perspective is eliding precisely the sort of deliberation about the nature of social welfare that justifies the administrative state in the first place.