

Artificially Intelligent Regulation

Mariano-Florentino Cuéllar & Aziz Z. Huq

This essay maps the potential, and risks, of artificially intelligent regulation: regulatory arrangements that use a complex computational algorithm or another artificial agent either to define a legal norm or to guide its implementation. The ubiquity of AI systems in modern organizations all but guarantees that regulators or the parties they regulate will make use of learning algorithms or novel techniques to analyze data in the process of defining, implementing, or complying with regulatory requirements. We offer an account of the possible benefits and harms of artificially intelligent regulation. Its mix of costs and rewards, we show, depend primarily on whether AI is deployed in ways aimed merely at shoring up existing hierarchies, or whether AI systems are embedded in and around legal frameworks carefully structured and evaluated to better our lives, environment, and future.

Unheralded and by inches, computational tools clustered under the label “artificial intelligence” are creeping into state and U.S. federal agencies’ toolkits for elucidating, implementing, and enforcing the law.¹ The Transportation Security Agency is required by law to deploy full-body millimeter-wave scanners trained to identify specific persons whose body shape indicates the need for further screening.² Sixty-three other civilian agencies of the federal government use more than 150 predictive tools to find facts, craft binding rules, exercise enforcement-related discretion, and detect violations of federal law.³ Local and state governments use similar tools to detect employment-benefit fraud, predict child abuse, and allocate police.⁴ In local criminal courts, prosecutors obtain convictions by drawing on probabilistic DNA analysis software.⁵ Local, state, and federal governments also leverage regulation to induce private parties to create and adopt new computational tools. The Department of Health and Human Services in 2016 created an algorithmic “security risk assessment tool” for health care providers needing to verify that their medical-records systems comport with federal data-security rules.⁶ Large investment banks increasingly adopt algorithmic tools as a means of complying with antifraud or money-laundering laws. Without fanfare, or wide public deliberation, the era of artificially intelligent regulation is almost certainly at hand.

We aim to map the potential, and risks, inherent in that new era. By *artificially intelligent regulation*, we mean regulatory arrangements that use a complex compu-

tational algorithm, or another artificial agent, either to define a legal norm or to guide its implementation. To see how AI might be integrated into the regulatory process at four distinct points – problem identification, empirical inquiry, rule formulation, and rule implementation – consider the following examples:

- A statute governing financial institutions' anti-money laundering responsibilities might define an explicitly Bayesian learning tool as part of an adequate anti-money laundering system. So long as a bank incorporates the tool, it would fall into a safe harbor against liability. The legally mandatory instrument, moreover, would dynamically update to account for new sorts of malfeasance at the regulatory authority's direction.
- Selecting particular people or families for the nation's refugee resettlement program, an agency might adopt as regulation a machine learning instrument to make acceptable decisions accounting for more vectors than can be easily calculated by a human decision-maker. The instrument will once again dynamically update to account for changing patterns of migration, geopolitical conditions, climatic change, and regional economic conditions. Human decision-makers might have to overcome a variety of challenges to take account of all of those relevant and complex streams of information quickly and accurately. An AI instrument could account for this information in a manner that contrasts with how human decision-makers would approach the problem without wholly breaking from the forms of human decision-making.
- A pollution emissions standard for manufacturing plants might be enacted as a reinforcement-learning algorithm. This instrument would define targets based on changing patterns of behavior and calculations of elasticity. It would hence respond dynamically to changing circumstances, including shifting strategies by emitting companies and their customers, quicker and cheaper than human modifications of regulation.
- An AI-infused mechanism might be adopted by law to allocate vaccines during a pandemic based on evolving data about a disease's spread, its symptoms, and the public's behavioral responses. Again, the regulation would take the form of a reinforcement-learning tool that changed based on evolving public-health circumstances.

These examples share common traits. Critically, in each one, *the law itself operates through a legally preordained computation process*. In the first, regulation defines compliance in terms of a (continually updated) algorithm. In the second, the distribution of state benefits is a product of an algorithm cast in law; distributions are not set in advance but emerge as a result of the algorithm's interaction with novel information. In the third, the algorithm-as-law defines a standard of conduct for private parties accounting for ways in which they, and others, respond in

real time. In the final example, the regulatory goal (defined as, say, maximum epidemic abatement) is formulated by hand in advance, but how that goal is realized is constantly recalibrated via computation of new data. In each of these use cases, a machine substantially displaces a different sort of human judgment.

The ubiquity of AI systems in modern organizations all but guarantees that regulators or the parties they oversee will make use of learning algorithms or novel techniques to analyze data in the process of defining, implementing, or complying with regulatory requirements. At one end of the continuum is the relatively incidental, isolated use of an AI system to assess whether data indicate that organizational enforcement priorities have changed over the course of a decade. At the other end of the spectrum is the statute that defines a financial institution's responsibility to guard against money laundering by formally defining, as a legal norm instantiated in a digital medium, a specific Bayesian updating function. Somewhere on the continuum one might draw a line to distinguish "artificially intelligent regulation" from more incidental use of manufactured intelligence merely to offer limited advice to legal decision-makers or evaluate the implementation of ordinary laws.

What to make of these arrangements is an intricate question that merits no simple answer. The public debate on regulatory AI is polarized between boosters and doomsayers. AI's diffusion across state instrumentalities hence provokes either shrill alarm or unblinking optimism. Minneapolis, San Francisco, and Oakland, for example, have all banned private facial recognition technologies that trawl public surveillance footage with AI tools. These jurisdictions enact the view that "AI is invariably designed to amplify the forms of power it has been deployed to optimize."⁷ In contrast, Chicago and Detroit recently purchased real-time facial recognition systems to integrate into their citywide camera networks.

We diverge from scholars who offer either pure celebration or lament about AI's effect on law. Rather, we think that artificially intelligent regulation holds promise and peril. As digitally native law, it exploits potential gains from new predictive technologies, and these gains have attendant costs and serious risks. We readily acknowledge AI's risks to human agency and democratic politics.⁸ We also think that the environmental impact of an industry already producing an estimated 3–3.6 percent of global greenhouse gas emission will also loom larger as usage increases.⁹ But we reject the broad claim that AI, as part of responsible social regulation with careful contingency planning and institutional safeguards, cannot deepen democracy, improve human welfare, or empower marginalized groups.¹⁰ Its mix of ensuing harms and rewards will instead depend on whether AI is deployed merely to shore up existing hierarchies, or whether its use aims to empower and better our lives, environment, and future.

We offer here an account of the possibilities of artificially intelligent regulation as a good and as a harm. We then offer thoughts on the "metaregulation"

of artificially intelligent regulation – that is, the larger regulatory frameworks in which agencies’ decisions to adopt or reject AI tools might be nested – within a democratic framework. Neither wholesale resistance nor an unthinking embrace of AI governance is justified. The national state and its agencies will almost certainly deepen entanglements with new predictive technologies. The ensuing form of artificially intelligent regulation, though, is not graven in stone. Experimentation with AI can help us better understand and resolve challenges arising from society’s often-conflicting expectations of the legal system for technical accuracy, democratic legitimacy, even-handed enforcement, and the nuanced consideration of situational factors. These various rule-of-law elements can be in tension with each other. AI systems can relax that tension, or perhaps exacerbate it in a specific case. But we see no alternative to the hard work of making sure that artificially intelligent regulation is designed to, and in fact does, advance the common good, and not deepen inequality or short-circuit democratic judgment.

Artificially intelligent regulation (AIR) is a legal norm that directly incorporates an algorithm capable of learning and adapting to new information, or the closely related activity of relying heavily on an algorithm to interpret or enforce a regulatory norm that may or may not itself directly incorporate an algorithm. The agency problem in regulation is familiar, but the AIR solution for it – and potentially achieving other goals – is novel.

We focus here on “regulation” in the sense of laws, rules, and guidance promulgated by an agency or department as part of an overarching legal framework for private activities like financial trading or health care. Regulation also includes the government’s efforts to control its own workings, such as policing and immigration. We do not address here the role of AI in the common law.¹¹ Our topic is distinct from discussions of “personalized” common-law rules of contract and tort law developed by courts rather than regulators.¹²

Our topic has analogies to certain long-standing arrangements in regulatory law. Some regulations already incorporate external standards by reference, such as industry norms, or encompass nontextual information.¹³ Although current administrative norms governing the *Federal Register* (the authoritative compendium of all regulations promulgated by agencies of the national government) may complicate the inclusion of a dynamic algorithm directly in a federal regulatory rule through incorporation by reference, both statutes and regulatory rules are sometimes drafted to allow agencies or the public to take account of changing knowledge or conditions.¹⁴ AIR can also act as a supplement or substitute for bureaucratically lodged discretion. The law is itself capable of evolving as agencies learn. Just as case-by-case adjudication elaborates the common law, so artificially intelligent law also adapts. But the locus of adaptation of AI is likely to be a standard internal to a statute or regulation, not a body of case law accreting over time.

Even well before legal norms become automated or intelligent, regulators will have little choice but to take seriously the world's increasing dependence on AI. The Internet shook governance beginning in the late twentieth century. It forced public agencies to contend, willingly or not, with new ways of disseminating information, networking computers, and shaping public perceptions.¹⁵ Regulators cannot unwind the widespread commercial adoption of AI techniques, such as backpropagation, neural nets, and large-language models, among contemporary firms. Algorithmic social media feeds, big-data trading platforms, and medical diagnostic tools powered by machine learning are, moreover, unlikely to be abandoned given consumer demand and the real welfare gains derived from them. Nor will regulated firms, including media platforms, banks, hospitals, and manufacturers, cease to innovate in respect to these tools – if nothing else because of unstinting foreign competition. The synergies between state and private enterprise in China, in particular, lend this commercial contest a geopolitical edge that cannot be wished away.¹⁶ Military agencies will keep pioneering technologies – like the communication protocols developed for the ARPANET project in the late 1960s that preceded the Internet – that invariably leak into civilian application. The conclusion that AI will increasingly infuse both government and society, therefore, is not mere lazy technological determinism. It is a reasonable inference from readily observable trends.

Still, invention is not the same as innovation.¹⁷ Not all digital tools catch on. The recent history of machine learning innovation has been uneven, punctuated by unexpected stops and starts. Whether new technologies are picked up, and how their benefits and costs are distributed, depends on social, economic, cultural, and even legal forces. However acute the pressures toward AI diffusion and adoption might be at this moment, nothing excuses regulators, jurists, and scholars from the difficult task of figuring out how those new tools are slotted into, and interact with, existing private or public institutions, as well as extant hierarchies coded by race, ethnicity, gender, or wealth. Nothing makes existing technological arrangements ineluctable. The monopolistic scale and network effects of dominant social media platforms, for example, was a contingent result of federal regulatory choices.¹⁸ Antitrust law might still find a way to reverse Facebook's and Google's dominance. Locally, the Los Angeles Police Department's April 2020 decision to abandon Palantir's crime-prediction software suggests that not all technological adoptions travel a one-way street. Predictions that AI inevitably serves to discriminate and disempower can enlist powerful historical examples. Their forward-looking force rests on a questionable disregard of democratic agency.

It would be a mistake to say that artificially intelligent regulation will ever completely displace human judgment in some form at some stage of the regulatory process. Human discernment designs and creates the learning tool embedded in the law. The fact that the application of rules to specific cases does not hap-

pen through the exercise of human discretion does not – indeed, cannot – lead to a complete absence of such discretion.¹⁹ There may not always be a human in (or on) the loop, but there is always a human exercising her judgment as to policy goals, what data are relevant to those goals, and how best to reconcile competing values: she may simply not be visible. With AIR, those judgments likely occur earlier in the design and implementation process. These judgments will tend not to be situated decisions, of the kind regulators now make, about how a norm applies to specific facts and particular persons. An instrument for matching refugees, for example, will not have information on particular flows of people, and almost certainly will be designed by engineers with little or no direct understanding of the refugee experience. As a result, the design stage of artificially intelligent regulations and the ensuing specification of predictive tools is a context in which biases (including invidious beliefs about race, gender, or other legally protected classifications), blind spots, and inaccurate generalizations filter into law. This human element of artificially intelligent regulation may well be occluded from the view of regulated parties.

Nevertheless, policy-makers and the public may have compelling reasons to move human judgment upstream and to filter it through a machine learning tool. At a very general level, AIR has the potential to make law and legal instruments more trustworthy – more amenable to accounting and discipline – and thereby to reduce the transaction costs of translating legal norms across different platforms and institutions.

The positive case for AIR comprises several elements. First, AIR can push agencies to define a societal goal more explicitly. Many AI instruments are organized around a “cost function” that examines each set of predictions of an outcome variable derived from historical data and defines a “cost” or penalty between predictions and the true (observed) outcome. The instrument is then trained to minimize that cost.²⁰ Writing a cost function requires a precise understanding of the social goals regulation seeks to advance. Because that judgment must be explicitly made, the cost function is an opportunity to air to the public both regulatory goals and the manner in which trade-offs are made.

A second benefit of AIR is flexibility over time. Agencies presently promulgate regulations and guidance as a means of implementing statutes periodically enacted by a legislature. Regulation often uses abstract or vague terms, or simply broadly sets a policy goal. Implementing that abstract statutory ambition – whether it is a safe workplace, a technologically feasible but environmentally tolerable level of emissions, or a decent refugee regime – requires translation. Regulators need then to write out their abstract goals in terms of particular rules or applications, bringing lofty aspiration into material form. AIR allows a well-informed legislature to install into law its abstract policy goal in a durable and adaptable way. Where regu-

lation adopts a reinforcement learning tool, the legislature also benefits from information that is not available at the time a law is passed. Hence, a resettlement algorithm might account for unanticipated shifts in migration patterns, or an antifraud tool could learn to recognize new species of criminal conduct. Thanks to this ability to build into law the capacity to dynamically update, a legislature condemned to only intermittent formal action via bicameralism and presentment is freed from frozen-at-the-moment-of-enactment text. This kind of flexibility may be especially valuable if the U.S. Supreme Court imposes new restrictions on Congress's ability to delegate through general grants of powers to federal agencies, with the latter filling in details with regulations.²¹ A law directing creation of an AIR might be a substitute for flexibility otherwise exercised through agency rulemaking over time.

A related benefit pertains to legislators' "agency cost" problem. Regulators may have different policy preferences from legislators. They might be excessively close to a regulated industry. Or they might slack off.²² One way to mitigate agency slack is with *ex post* judicial review. But the use of courts as an oversight mechanism has costs. Litigation can be used to delay desirable regulation. Fearing a suit, budget-constrained agencies might forgo action. Regulated parties, anticipating judicial review, have an incentive to lobby for particular judicial appointments.²³

AIR addresses agency slack in a different way. By impounding their judgments into a digitally native tool, legislators drain away later discretion about how a law is enforced. The resources used up in translating verbal standards back and forth to code and mathematical specification are likely to be smaller than the social resources sucked up by litigation clashes between interest groups and the government. AIR, however, does not eliminate agency problems entirely. Realistically, legislators must rely on technologists and coders to craft an instrument. Unless a legislator can trust the designers of digitally native law, as well as the sources of training data, the specter of "capture" and distorted preferences arises once more.²⁴ Legislators could demand benchmarking and transparency in AI design "appropriate for practical reasoning," not just in terms of technical detail.²⁵ Such arrangements might further facilitate either *ex post* judicial review (especially when individual rights are at issue) or legislative committee-based oversight.

The advance of AIR under these conditions may also alter democratic governance more broadly. At present, a legislature enacts a law with limited control over how its terms are understood and applied in the future. Later legislatures can use their appropriations power and their ability to jawbone agency leaders to nudge regulation toward their preferred policies, even when they diverge from those of the enacting legislature. Sometimes, courts step in to interpret statutes in ways that force the later legislature to act more overtly by passing new law.²⁶ But not always, and not reliably.

AIR might scramble such arrangements. In principle, it empowers an enacting legislature. That body has the ability to enact not just the law in an abstract form,

but also to embed a mechanism for updating. This sharpens the importance of the discrete political moment in which a law is enacted; it also diminishes the importance of the legislative power to influence agencies in the long term. Arguably, this is salutary in terms of democratic norms. It ups the stakes of the actual legislating moment, when the voting public is most likely attentive, while diminishing the importance of periods in which the public is less engaged, and legislative influence more diffuse. This helps voters exercise retrospective judgment about their representatives. On the other hand, AIR, in a paradoxical sense, by making formal laws less brittle and more capable of built-in adaptation, could conceivably enable long-past legislative coalitions to endure beyond their expiration date. Hence, it may empower the dead hand of the past against the influence of living legislators wielding a current democratic imprimatur.

Finally, it is worth considering whether AIR can be used to broaden access to legal institutions and the benefits of law more generally. Algorithmic tools already facilitate estate planning via websites such as LegalZoom. While these instruments are not without complications, it is worth considering ways in which AIR might be used to empower ordinary citizens presently discouraged from seeking legal remedies by litigation's complexity and cost.²⁷ This is one important way of resisting the complacent assumption that AI is an innovation that necessarily and inevitably concentrates power and increases pernicious social inequalities.

All should not be presumed to be well with this potential new era of regulation. Just as it enables optimal adaptability, diminished agency costs, and lower transaction costs, so artificially intelligent regulation will engender new problems of transparency, legitimacy, and even equity. All raise fundamental questions of constitutional magnitude.

To begin, it is premature to assume AIR always reproduces undesirable or malign forms of hierarchy. Though regulation is not guaranteed to enhance social welfare, neither is it intrinsically regressive. It has advanced the cause of civil rights, workplace safety and health, environmental protection, and consumer rights. AIR is just one species of regulation. Of course, all lawmaking risks interest-group capture or the unintended perpetuation of invidious stereotypes. AIR, like any kind of legal intervention, must be scrutinized for those risks. In particular, AIR empowers a new class of experts – computer scientists and engineers – at present noteworthy for its lack of gender, racial, and ethnic diversity. Finding diversity in such expertise and turning the latter to serve the public good is not impossible: biological and medical science has shown as much. But it will require sustained institutional change.

More seriously, the ends and means of AIR – like many of the complex statutes that Congress, in particular, has enacted – are not necessarily readily perceived or understood by nonspecialist members of the public or elected officials. The value

of an algorithm that diminishes conventional principal-agent problems involving human-led agencies also means that standards might evolve in problematic ways. This may be a result, for example, of mistakes in how a reinforcement learning reward function is specified, or it can be a consequence of adversarial disruptions. There is a question whether any adaptational “drift” distorts what the law achieves, or instead demands fresh involvement by the very mix of experts, politically accountable officials, and competing stakeholder pressures that optimistic proponents might expect these new forms of law to render unnecessary.

A yet more fundamental question is whether an AI-based legal arrangement would be perceived as legitimate in either a sociological or legalistic sense. The ability of the public to understand what AI does at the front end is limited, although that is also true of many existing laws and legal institutions. Leaving aside the precision with which a dynamic legal provision “aligns” with a defensible macroconcept of social welfare in advance, the way such provision *evolves* over time is not made legitimate without further ingredients. These include the capacity of concentrically larger circles of people, including agency officials and regulated parties at minimum, to understand certain things about how a system performs. Also relevant is affected parties’ capacity to argue in terms the public understands about why AIR is performing adequately (or not) relative to the rest of the jurisdiction’s legal commitments. Agencies or lawmakers could also create “tripwires” to prevent excesses in the use of public, coercive authority; capture or co-optation by private interests through de facto private delegation; or violations of due process, equal protection, or anticorruption norms. Certain uses of coercion may also be *ipso facto* illegitimate without human oversight.

Such measures could be calibrated to promote institutions that allow debate about how a law gets implemented in a particular situation and about the policy and value assumptions supporting the law. Equally important are arrangements that prevent the use of AIR as a shield to prevent public accountability for the coercive use of power. Here, “public accountability” means that some people must accept responsibility for the use of coercive authority in ways that account for material and emotional consequences, including loss of income, reputational degradation, loss of interesting work, and misrecognition by peers or authority figures.

Finally, with opacity comes the risk that algorithms reenact malign hierarchies of race, ethnicity, class, and gender via inscrutable code and invisible design choices.²⁸ The terminology of “bias” in AI is used in varied and inconsistent ways. In our view, the most powerful normative concerns arise when the use of AI imposes material harms on a historically subordinated group.²⁹ Machine bias defined in this way need not flow from any conscious decision to suppress a historically subordinate group. It can result simply from inattention or ignorance by programmers who are not members of those classes. Preventing intentional or inadvertent reproduction of these hierarchies requires active attention to the code

inserted into regulation. As recent turmoil at Google’s ethics division suggests, the implementation of equity is no simple matter, but demands organizational leadership and effective staffing.

None of this means AIR should be eschewed. But technical limitations and public resistance mean AIR will likely be limited in scope for some time. More interesting to us is how the emergence of AIR raises questions about the “metaregulatory” structure of administrative and regulatory law. That is, how should the law itself guide the creation and oversight of digitally native law?

The law has already developed tools to audit and evaluate ordinary regulation; cost-benefit analysis is foremost among them. AIR requires rethinking and retooling government’s auditing and oversight capacity to extend values of equity and rationality into its frameworks. As governments create “sandboxes” in which to build and test AIR, they will need to apply robust norms of transparency and benchmarking to ensure that AI is not just the product of – but also facilitates – reasonable and informed deliberation. Experiments with AIR may benefit from building in some of the encumbrances that surely make laypeople wonder about law as it operates today. Digitally infused regulations might therefore explicitly incorporate interpretive mechanisms that will “translate” a standard into ordinary language. Periodic audits for practical bias along race, gender, and other lines might be mandated by law, with failure to pass connected to a penalty of statutory rescission.

More broadly, a federal agency can be imagined for being responsible for sourcing, testing, and auditing new digital tools. Such an agency could be especially helpful given some of the persistent difficulties public organizations face when making procurement decisions. The agency would benefit from a capacity to experiment with recruitment and retention tools, including rotation and part-time arrangements subject to appropriate safeguards against conflicts of interest, to layer into the highest levels of the public sector the kind of expertise and mix of cultures helpful in enhancing government capacity for assessment of AIR. It would operate much as the General Services Administration, established by President Harry Truman in 1949, serving as a hub for digitally native law, a source of auditing expertise, and a locus for public complaint.

For the foreseeable future, AIR offers fascinating possibilities for enhancing governance, but it will nonetheless face intense constraints. Given the risks entailed, perhaps this is as it should be. If AIR is to become legitimate, it must face a trial by fire under the abiding rule-of-law constraints familiar from our existing, imperfect legal system. Further, it will be subject to the coterie of pluralistic pressures capable of creating such enormous friction for even the most elegantly designed legal reforms. Both will confer legitimacy and limit risks of severe

error, but also erode AIR's possibilities and promise. Perhaps such friction is not entirely useless. Perhaps, indeed, it has the potential to force nuance into discussions about how to reconcile contending ideas about what sort of social welfare regulation is supposed to advance. The resulting constraints also offer a powerful reminder that the social benefits of AIR depend at least as much on our society's capacity to engage in intelligent governance as they do on continued progress in machine learning.

ABOUT THE AUTHORS

Mariano-Florentino Cuéllar, a Fellow of the American Academy since 2018, is President of the Carnegie Endowment for International Peace. A former Justice of the Supreme Court of California, he served in the Obama administration as Special Assistant to the President for Justice and Regulatory Policy at the White House Domestic Policy Council. He was previously the Stanley Morrison Professor of Law at Stanford University and is the author of *Governing Security: The Hidden Origins of American Security Agencies* (2013).

Aziz Z. Huq is the Frank and Bernice J. Greenberg Professor of Law at the University of Chicago Law School. He is the author of *The Collapse of Constitutional Remedies* (2021), *How to Save a Constitutional Democracy* (with Tom Ginsburg, 2018), and *Unbalanced: Presidential Power in a Time of Terror* (with Frederick A. Schwarz, 2007).

ENDNOTES

- ¹ 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.: Administrative Conference of the United States, 2020).
- ² *Federal Register* 81 (42) (2016): 11363.
- ³ Engstrom et al., *Government by Algorithm*, 15–16.
- ⁴ Virginia Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor* (New York: Macmillan, 2018).
- ⁵ *People v. Chubbs*, 2015 WL 139069 (January 9, 2015).
- ⁶ “Security Risk Assessment Tool,” Health IT, <https://www.healthit.gov/topic/privacy-security-and-hipaa/security-risk-assessment-tool>.
- ⁷ Kate Crawford, *Atlas of AI* (New Haven, Conn.: Yale University Press, 2021), 224.
- ⁸ Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: Public Affairs, 2019).

- ⁹ Lotfi Belkhir and Ahmed Elmehri, “Assessing ICT Global Emissions Footprint: Trends to 2040 & Recommendations,” *Journal of Cleaner Production* 177 (1) (2018): 448–463.
- ¹⁰ Crawford, *Atlas of AI*, 223.
- ¹¹ Mariano-Florentino Cuéllar, “A Common Law for the Age of Artificial Intelligence,” *Columbia Law Review* 119 (7) (2019): 1773–1792.
- ¹² Ariel Porat and Lior Jacob Strahilevitz, “Personalizing Default Rules and Disclosure with Big Data,” *Michigan Law Review* 112 (8) (2014): 1417–1478.
- ¹³ Federal regulations may incorporate “published data, criteria, standards, specifications, techniques, illustrations, or similar material.” 1 CFR § 51.7 (a)(2)(i); and Emily S. Bremer, “On the Cost of Private Standards in Public Law,” *Kansas Law Review* 63 (2015): 279, 296.
- ¹⁴ Office of the Federal Register, *IBR Handbook* (Washington, D.C.: Office of the Federal Register, 2018). See, for example, *Whitman v. American Trucking Associations*, 531 U.S. 457 (2001), discussing the EPA’s responsibility under Section 109(b)(1) of the Clean Air Act to set National Ambient Air Quality Standards that are “requisite to protect the public health.”
- ¹⁵ See Jonathan Zittrain, *The Future of the Internet and How to Stop It* (New Haven, Conn.: Yale University Press, 2008).
- ¹⁶ Aziz Z. Huq and Mariano-Florentino Cuéllar, “Privacy’s Political Economy and the State of Machine Learning,” *NYU Annual Survey of American Law* (forthcoming).
- ¹⁷ David Edgerton, *The Shock of the Old: Technology and Global History Since 1900* (London: Profile Books, 2011).
- ¹⁸ Dina Srinivasan, “The Antitrust Case against Facebook: A Monopolist’s Journey towards Pervasive Surveillance in Spite of Consumers’ Preference for Privacy,” *Berkeley Business Law Journal* 16 (1) (2019): 39–99.
- ¹⁹ Aziz Z. Huq, “A Right to a Human Decision,” *Virginia Law Review* 106 (3) (2020): 611–688.
- ²⁰ David Lehr and Paul Ohm, “Playing with the Data: What Legal Scholars Should Learn about Machine Learning,” *UC Davis Law Review* 51 (2) (2017): 653–717.
- ²¹ Cass R. Sunstein, “The American Nondelegation Doctrine,” *George Washington Law Review* 86 (5) (2018): 1181–1208.
- ²² Kenneth J. Meier and George A. Krause, “The Scientific Study of Bureaucracy: An Overview,” in *Politics, Policy, and Organizations: Frontiers in the Scientific Study of Bureaucracy*, ed. George A. Krause and Kenneth J. Meier (Ann Arbor: University of Michigan Press, 2003): 1–19.
- ²³ Cass R. Sunstein, “On the Costs and Benefits of Aggressive Judicial Review of Agency Action,” *Duke Law Journal* 1989 (3) (1989): 522–537.
- ²⁴ Justin Rex, “Anatomy of Agency Capture: An Organizational Typology for Diagnosing and Remediating Capture,” *Regulation & Governance* 14 (2) (2020): 271–294.
- ²⁵ John Zerilli, John Danaher, James Maclaurin, et al., *A Citizen’s Guide to Artificial Intelligence* (Cambridge, Mass.: MIT Press, 2021), 33.
- ²⁶ Einer Elhauge, “Preference-Eliciting Statutory Default Rules,” *Columbia Law Review* 102 (8) (2002): 2162–2290.

- ²⁷ Emily S. Taylor-Poppe, “The Future Is Complicated: AI, Apps & Access to Justice,” *Oklahoma Law Review* 72 (1) (2019): 185–212.
- ²⁸ Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York: NYU Press, 2018).
- ²⁹ Aziz Z. Huq, “Racial Equity in Algorithmic Criminal Justice,” *Duke Law Journal* 68 (6) (2019): 1093–1134.