DATA AND DEMOCRACY



Is the Administrative State Ready for Big Data?

By Wendy Wagner and Martin Murillo

KNIGHT
FIRST AMENDMENT
INSTITUTE at
COLUMBIA UNIVERSITY

In October 2020, the Knight First Amendment Institute at Columbia University convened a virtual symposium, titled "Data and Democracy," to investigate how technological advances relating to the collection, analysis, and manipulation of data are affecting democratic processes, and how the law must adapt to ensure the conditions for self-government. This symposium was organized by the Institute's 2019-2020 Senior Visiting Research Scholar, Yale Law Professor Amy Kapczynski, and co-sponsored by the Law and Political Economy Project at Yale Law School.

The essays in this series were originally presented and discussed at this two-day event. Written by scholars and experts in law, computer science, information studies, political science, and other disciplines, the essays focus on three areas that are both central to democratic governance and directly affected by advancing technologies and ever-increasing data collection: 1) public opinion formation and access to information; 2) the formation and exercise of public power; and 3) the political economy of data.

The symposium was conceptualized by Knight Institute staff, including Jameel Jaffer, Executive Director; Katy Glenn Bass, Research Director; Amy Kapczynski, Senior Visiting Research Scholar; Alex Abdo, Litigation Director; and Larry Siems, Chief of Staff. The essay series was edited by Glenn Bass with additional support from Lorraine Kenny, Communications Director; A. Adam Glenn, Writer/Editor; and Madeline Wood, Communications and Research Coordinator.

The full series is available at knightcolumbia.org/research/

INTRODUCTION

LAWS THAT PROTECT the environment and public health are notoriously dependent on data-intensive analyses to inform the setting of protective standards, effluent limits, licensing decisions, and enforcement priorities. Thus, technological innovations in sharing, collating, and analyzing data—particularly in the wake of advanced forms of artificial intelligence (AI)—hold the promise of revolutionizing this sector of administrative law.[1] Regulatory projections that in the past could only be done, at best, by hand in resource-intensive ways can now be completed quickly and cost-effectively using algorithmic tools. For example, machine-learning models can make predictions on the toxicity of individual, "mystery" chemicals by processing all available data on related chemical structures to provide a predictive portrait of the potential risks of these new chemicals.[2] Existing data on air, water, and soil pollution can be assembled to yield community assessments of past and predicted future pollutant loads that direct regulator attention to neighborhood hot spots, inequities, and even potential facility violations of existing legal requirements.[3] Product manufacturers' mandated submissions of one-off, "adverse effects" occurring post-market can be synthesized by sophisticated computational tools to identify those products that warrant

immediate attention by regulators.[4] And that is only the beginning.

Yet, is the administrative state truly ready to make use of these new, sophisticated computational tools in accountable and rigorous ways? At first blush, the answer would seem to be a resounding "yes." Given administrative law's foundational commitment to deliberation, reason-giving and expertise, this particular part of our legal system seems well-positioned to make excellent use of these tools for the public good.

However, a deeper look at the scaffolding of administrative requirements and the resulting incentives reveal some areas of concern. In this article, we provide a critical analysis of administrative law as it applies to "algorithmic tools," and locate some pockets of significant misalignment between administrative law and the effective, accountable and even innovative use of algorithmic tools. Specifically, we identify a potentially wide gap between what agencies *should* do based on elemental best practices for the accountable use of algorithmic tools, and what agencies are *likely to do* under the directives of administrative law.[5]

But the incentives implicit in administrative law may be even more problematic. Rather than encourage and reward agencies that employ best practices to ensure the accountable development and use of algorithmic tools, in some regulatory sectors, administrative law sometimes tacitly rewards agencies for developing and using algorithmic tools that are opaque and potentially biased. We argue that, at the very least, the guiding hand of administrative law is not pointing agencies in the right direction.

Our paper unfolds in three parts. In the first part, we collect some emerging, consensual principles on best practices for the publicly accountable development and use of algorithmic tools for informing policy. In the second section, we explore administrative process using a kind of troubleshooting approach to assess how well its legal requirements map against these best practices, particularly when algorithmic tools are used to inform agency decisions that are judicially reviewable. We identify a number of counterproductive incentives emerging from this mapping exercise. In the third section, we offer recommendations for reform.

One important qualification bears mention before proceeding. Since the administrative state is vast and tends to resist generalization, we focus on the use of algorithmic tools in public health and environmental regulation, a

focus that is justified in part because it is an important area of social regulation for which advanced computational tools are expected to be particularly beneficial.[6] Within this topic area, we then narrow our focus still further by examining agency use of algorithmic tools that inform or create binding rules subject to notice and comment, and judicial review. The added scrutiny of stakeholders and courts under the Administrative Procedure Act (APA) would seem to present one of the best cases for judging how well the law is currently situated to encourage the accountable use of these technologies.[7]

I. ALGORITHMIC TOOLS: ACCOUNTABILITY RISKS AND BEST PRACTICES

At least since the Cold War, modeling and algorithms have been instruments used by regulatory and government decision-makers in a variety of beneficial (and not so beneficial) ways.[8] More than a half-century ago, space and military agencies began to utilize mechanistic and statistical models—often implemented through algorithms—to optimize trajectories, fuel, energy, and other variables.[9]

While the use of algorithmic tools to support governmental activities have been steadily increasing for decades, significant advances in data science, particularly the emergence of advanced artificial intelligence/machine learning (AI/ML) algorithms, such as deep learning or natural language processing, along with exponentially more data and processing power is becoming transformative. However, unlike mechanistic models that describe the inputs and outputs of a phenomena using rules developed by humans with an in-depth knowledge of the phenomena in question, the application of advanced forms of AI/ML generate the rules "on their own" from large amounts of data.[10] As a result, these AI/ML models "lack explainability at a fundamental level," presenting serious accountability challenges when used for policy.[11]

Given important differences between the more "inscrutable" advanced AI algorithms and older, more intuitive models and associated algorithms, many researchers have focused their sights exclusively on the intersection between these newer AI/ML technologies and the law. But in at least some

parts of the administrative state, the core challenges facing the new wave of AI/ML are different, albeit only in degree, from the long-standing accountability challenges posed by complicated mechanistic models and associated algorithms. Because all of these tools offer algorithmic representations meant to approximate reality, most members of the public, their representatives, and even policymakers cannot make sense of their results without considerable education and assistance.[12]

To ensure that we do not slice the problem too narrowly, for the purposes of this paper we investigate the use of what we call "algorithmic tools" more generally, focusing mainly on advanced forms of AI, but also including elemental forms of ML, mechanistic models, and all algorithms that are used to support and inform policy.

A. Accountability challenges for algorithmic tools

One of the primary concerns about algorithmic tools, old or new, is the challenge they create for democratic processes. Without rigorous public and expert oversight, algorithmic tools can obscure important decisions and discretionary choices from public view, introduce bias, be rife with errors, or ultimately constitute Trojan horses that make it possible for privileged stakeholders to hijack the administrative process. The lack of multidimensional and interdisciplinary mechanisms to guide the development and implementation of these tools will not only inhibit early detection of these risks, but could invisibly undermine democratic decision processes themselves.

Yet meaningful public involvement in overseeing the development of these tools is challenging, due to the technical difficulties involved in evaluating the tools and their outputs, the lack of in-house expertise, and the absence of enabling structures necessary to participate effectively. Affected stakeholders, particularly thinly financed groups attempting to protect the interests of the broader public, may run up against a "computational wall" that prevents them from penetrating the algorithms and underlying policy choices incorporated into a tool developed by an inner circle of well-financed stakeholders and agency staff. But even some insiders may find themselves lost in the source code, unable to extract even the most basic questions the modeling process was designed to answer.[13]

If anything, these concerns are magnified for AI/ML tools because

advanced AI algorithms are inherently "black boxes" that discern patterns and make predictions in ways that cannot be intuitively understood or explained in the same way as more basic algorithmic tools.[14] Moreover, in contrast to mechanistic or statistical models, advanced forms of AI/ML typically include a new, active group of experts who are more sophisticated in applying data science techniques but who can be inexperienced in the phenomena in question. This is further complicated by the fact that there exists a wealth of discretionary choices that need to be made at all stages of the algorithmic tool's development.

Scholars have addressed some of these concerns through the principle of algorithmic accountability.[15] Since participation is a core component of administrative decision-making, a constant lodestar for the use of these algorithmic tools is that the high tech tools *must* be used in ways that are accountable to the public.[16] Yet, to be meaningful, democratic accountability in the use of algorithmic tools requires not only public direction and engagement, but also unbiased expert review to assess the technical integrity and policy fidelity of the tools on behalf of the public. This dual oversight by both experts and the general public is vital to ensuring that algorithmic tools are developed in ways that minimize the aforementioned risks and advance the broader public interest.[17]

Two rather obvious challenges arise for agencies in meeting this accountability goal for algorithmic tools. First, in contrast to the default assumption that "sunshine is the best disinfectant," there is general agreement among commentators that, with respect to most algorithmic tools, it is not enough to simply provide a kind of "fishbowl" transparency[18] or open access to effectuate meaningful public engagement.[19] Rather, given the high costs of making sense of the technical tools, a much more elaborate translation effort is needed to solicit needed guidance and feedback from a range of different participants in the development and use of algorithmic tools, including expert review.[20]

Second and relatedly, because the algorithmic tools are complex, there are likely to be more obstacles to ensuring that this engagement includes the *full range* of affected parties (experts and public). Diverse scrutiny is vital to science[21] and is also vital to public policy. Given the many judgments embedded in algorithmic tools, it is essential that *all* affected public stakeholders

and diverse experts (and not simply a small subset of well-financed stakeholders) be involved in the identification and development of these tools on behalf of the public.[22] Only by leveraging the participation of diverse stakeholders can these otherwise inscrutable tools be developed and evaluated in rigorous ways that fulfill fundamental principles of accountability.

B. Best practices to enhance accountability

So, from high altitude, what practices are needed to guide the design, development, and utilization of algorithmic tools so that they may be used accountably?[23] The actual design of the ideal institutional architecture needed to support diverse and meaningful engagement with experts and the public, not surprisingly, has been a work in progress for decades.[24] But the emergence of advanced forms of AI/ML, their data requirements, and the proprietary nature of some of the resources used in their development makes the challenges inherent in this endeavor more acute.

By synthesizing the literature, including early recommendations for ensuring the accountability of mechanistic and statistical models by the National Academies of Science, we endeavor to extract some preliminary "best practices" in order to establish a multidimensional framework for the development and use of algorithmic tools that provide a kind of "accountability by design"[25] (see Figure 1). At base, these best practices insist on structured and carefully explicated steps throughout the tool's development and use, which help ensure meaningful engagement with experts and the public from "the start."[26] This explication also spotlights underlying substantive policy choices to advance public engagement.[27]

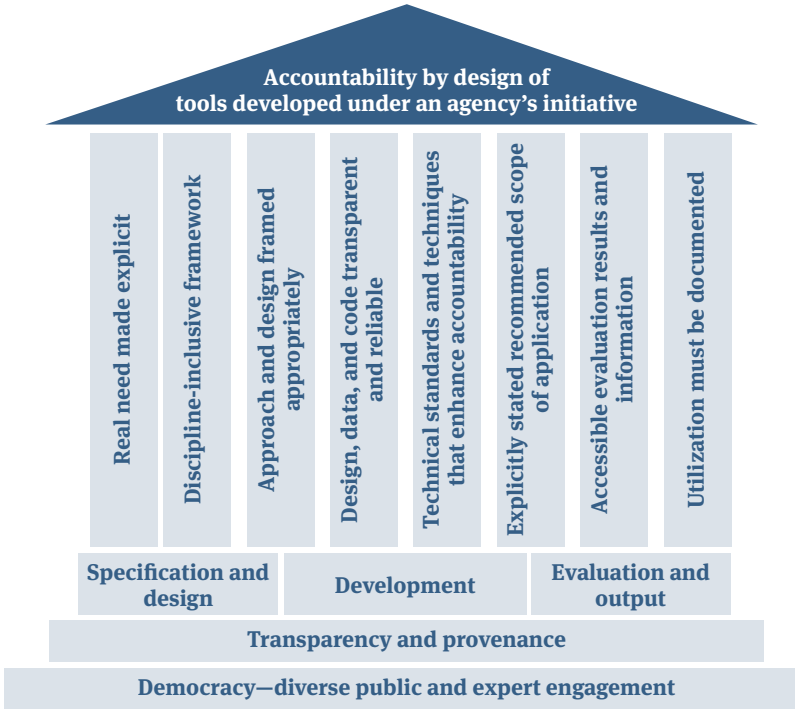


FIGURE 1: REPRESENTATION OF A FRAMEWORK FOR DEVELOPING ACCOUNTABLE ALGORITHMIC TOOLS THAT REST ON GENERAL BEST-PRACTICE PRINCIPLES

Best practices not only insist on fishbowl transparency and the clear explication of choices throughout the development of algorithmic tools, but also the comprehensive documentation of how a model was assembled, and the roles played by authors and participants throughout development of the tool.[28] Since model development involves so many different players, an elaborate provenance statement (how it was developed, by whom, the levels of review, and the conflicts of interest at each step)[29] is vital to allow various public and expert audiences to assess the integrity of the process, as well as ensure adequate deliberation and oversight at key steps in the modeling process.

These best practices also require that the needed explications and choices reach two audiences—expert data scientists, and policymakers and

the public.[30] Excellent data science practices will not suffice to ensure the accountable use of algorithmic tools. For example, an impeccably developed AI model from a data scientist's perspective might still fail best practices if the catalog of key choices (e.g., the purpose of the model and its generalizability, how it will be evaluated) are not explained in ways that link the model to policy goals and policy consequences. Conversely, a clearly explained policy purpose for an AI model might still lack needed direction for data scientists, leading to problematic drift in how the model is developed and used. Explications are essential, but they are also tricky given the multiple audiences occupying this thick interdisciplinary space.

1. SPECIFICATION AND DESIGN OF THE TOOL

Real need made explicit: Since algorithmic tools seek to answer a question or solve a problem in the policymaking process, it is vital that a proposed tool's purpose, goals, and decision criteria be identified explicitly at the start of its development.[31] Thus a clear statement of the policy purpose and concept of the algorithmic tool should be accessible to a variety of affected groups in order to provide a meaningful democratic deliberative process.[32] In the absence of such a public statement, the algorithmic tool's framing and purpose becomes more malleable and open-ended in ways that create the opportunity for significant sources of bias and counterproductive "tweaking" to creep in.[33]

Discipline-inclusive framework: The development of an algorithmic tool must also be open to alternative disciplinary approaches. The risk associated with a lack of inclusiveness is that a small set of stakeholders or experts, well-versed in a single approach to addressing the issue, might monopolize the process, thus inhibiting the work of other disciplines that have already addressed similar issues.

2. DEVELOPMENT OF THE TOOL

Approach and design framed appropriately: During the development process, the selection of algorithms and data will be based on a range of sources that can include raw policy preferences, scientific theories, empirically identified associations, and statistical computations.[34] For our analysis here, it

is only important for readers to appreciate the best practices the agency is expected to use to explain and justify these numerous consequential decisions for scientists, the public, and lay policymakers. While it becomes difficult to explicitly detail the algorithmic logic that underlies AI, particularly for advanced forms of AI, emerging fields such as explainable AI can be used to lay bare a model's decision "pathway," using such tools as visualizations and ranking features based on their pivotalness in the model.[35] Standards organizations such as the National Institute of Standards and Technology (NIST), the Institute of Electrical and Electronics Engineers (IEEE), the International Standards Organization (ISO), and many others provide important guidelines and certifications for making software technically and universally "accountable" for the benefit of society.

Code and data sets must be transparent: Although fishbowl transparency is not sufficient for the accountable use of algorithmic tools, it is nevertheless necessary. The underlying purpose and specification statements, source codes, data sets, user guides, and other information should all be easily accessible for both the public and experts alike in order to allow them to reproduce outputs, test individual modules, and compare or benchmark the tool against other possible solutions or outputs[36]. To ensure this "library" of publicly available information is useful in practice, agencies must make a distinction between what is essential material and what is non-essential material (and identify what different types of audiences might consider essential). Note that when source codes and data sets are proprietary or confidential, it will be even more critical to focus attention on comprehensive provenance statements and reasoned transparency.[37]

Underlying (i.e., training) data must be reliable: The limitations of data, including the training data, must be carefully detailed, and the justifications for specific data choices made clear. Biased or distorted training data will obviously bias the resulting model.[38] When privacy or other forms of confidentiality may limit the training data that is publicly available, creative techniques must be developed to ensure adequate public and expert oversight. For example, sampled or anonymized versions of the data can be provided in addition to results of statistical analyses and visualizations. Furthermore,

when fishbowl transparency is not possible, provenance statements and reasoned transparency must be made even more elaborate and perhaps even certified by oversight bodies with full access to code and data.

Recommended scope of application: The range of applications for which the algorithmic tool was originally created must be carefully and accessibly explained. This explication of the generalizability of the tool determines the range of issues to which it reliably applies.[39] For example, if the training data only considers a subset of the population, then this limitation should be considered in determining how the tool is used for policy and comprehensively documented.

3. EVALUATION METHODS AND OUTPUT

Accessible evaluation results: The process for evaluating the algorithmic tool—which could take many forms—is a vital feature of its development that "continues throughout [its] life."[40] Evaluation processes and results, as well as the underlying algorithms and data sets, must be made transparent so that stakeholders can benchmark outputs against other alternatives, including currently applied mainstream processes. As Kroll et al. point out, "Transparency and after-the-fact auditing can only go so far in preventing undesired results."[41] A lengthy chapter on evaluation of models in the National Academies' Models report recommends that evaluation of algorithmic tools include both ex ante and ex post commitments that bind the evaluation process in advance.[42]

Documented utilization: The specification of a real need and the eventual utilization of the tool are first and last stages of the best practices. Accordingly, the public (and expert peers) should be able to locate and trace how the tool is used to affect regulatory outcomes and determine whether those uses are appropriate. This is particularly important if the tool's role in the regulatory process is not well defined, thus potentially compromising meaningful public input and oversight.

A number of algorithmic tools are currently developed in the "sandbox" or behind closed doors; some of these trial-and-error models may ultimately

"slip in" to inform policy. It is critical that all algorithmic tools be subject to a rigorous evaluation that includes transparency and provenance requirements, as well as a rigorous evaluation and an identification of how these tools meet the best practices.[43]

II. ADMINISTRATIVE PROCESS AND ALGORITHMIC TOOLS: A CRITICAL ANALYSIS

GIVEN THIS VITAL ROLE of best practices to guide the use of algorithmic tools in agency decisions, how does administrative law encourage agencies to make the best use of them, including for advanced forms of AI? After offering orientation to administrative law, we discuss how these principles fit within the demands of administrative law, with a primary focus on algorithmic tools that inform or support agency rules. This could include, for example, complex risk assessment models used to estimate safe levels of a pollutant or AI-generated predictions of toxicity that lead to regulatory restrictions. (Indeed, the tool may even constitute a rule itself.)

We find that while existing law allows agencies to use AI and related algorithmic tools to inform regulatory decisions,[44] the law does a poor job of guiding the agency's development and incorporation of these algorithmic tools into policymaking in constructive ways.[45] We thus join a growing chorus of scholars who observe a "substantial mismatch" between existing administrative process and "the technical demands of algorithmic oversight."[46] At the same time, while we do not mean to understate the particularly difficult "explainability" problem at the intersection of AI/ML and the law,[47] our analysis reveals that this is not a wholly new problem, but rather a more difficult permutation of an existing, long-standing challenge to the agencies' use of complex technical evidence.[48]

This section begins with a quick refresher on how administrative process works. We then discuss how existing legal oversight and accountability mechanisms not only fail to require agencies to detail out the best practice structures, but serve to impose a series of deterrents that could even undercut agency efforts to use algorithmic tools in accountable ways that advance public-benefitting goals.

A. Primer on administrative process

The administrative state is premised on a set of processes that seek to hold agencies accountable to the public, despite the fact that agency officials are unelected. Two features—democratic deliberations and reason-giving—are particularly critical to ensuring the legitimacy of agency decisions.[49]

The Administrative Procedure Act of 1946 establishes the basic framework for inculcating these democratic goals into agency decisions.[50] Under the APA, agencies' promulgation of binding rules requires the agency to publish a notice of proposed rulemaking in the Federal Register and, at a minimum, solicit written comment from the public concerning the merits of its proposal. Commenters who find their submitted concerns insufficiently addressed can then challenge the rule in court, arguing that the rule is "arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with law."[51]

Public participation in notice and comment thus plays a fundamental role in ensuring the accountability of agency rulemakings and the underlying analyses (see Figure 2).[52] Participants can not only challenge agencies in court but can send fire alarms to elected officials in Congress and the White House encouraging them to intervene. Additionally, participants can, and often do, play additional, more informal roles in providing key information and resources (such as models) to inform the agency's formulation of the proposed rule.[53]

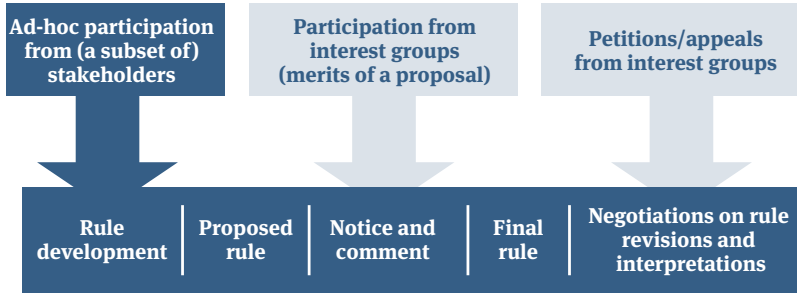


Figure 2: The stages of participation in agency rulemakings

As a result of this procedural design, agencies are typically careful to ensure that they have adequately anticipated and responded to critical substantive comments relating to their analyses and rules.[54] Agencies also generally take care in the final rule to provide the "reasoning" that supports contested features, since a failure to provide this "reasoned decision-making" can lead the rule to be reversed and remanded.[55] While courts tend to be quite deferential to the agency's explanations of its ultimate decision, some courts do take a "hard look" to ensure the responses are adequate.[56]

Because the role of public commenters is so important to the substance of the rule, and also because much of the information the agency needs to fulfill its mandate rests in the private sector—often with the industries it regulates, administrative deliberation tends to be open-ended and dynamic.[57] There are generally no requirements that restrict the agency from working closely with stakeholders, including a subset of stakeholders, in developing the rule proposal.

Finally, under existing administrative process, political appointees or even the Office of the President can shape the rules (as well as other agency policies), even if they are based on highly technical evidence and analysis.[58] While the political influences likely apply only to the most significant rules and policies, much of the political intervention occurring during the analysis and rulemaking processes is invisible.[59] Criticisms of the Trump administration, for example, detail very specific examples of the White House reaching down into the technical features of the analyses supporting rules and policy decisions.[60]

B. Areas of slippage and misalignment between the law and best practices

Process requirements of "democratic deliberations" and "reason-giving" seem well-suited, in the abstract, to encouraging the rigorous and accountable use of algorithmic tools for rulemakings. Ideally, the agency will include in their "reasons" structured information about the analytical steps (e.g., framing, choices, evaluation) and the provenance of the tool (e.g., model). Ideally, too, democratic deliberations will involve a full mix of participants at each of the key stages of a tool's development to ensure that the agency is accessing all of the relevant information and subjecting its choices to rigorous and diverse oversight.

But, as we detail below, not only is the "ideal" not required in practice, there are reasons to worry that, as currently designed, administrative processes might actively discourage agencies from engaging in best practices in order to survive and navigate the adversarial and often convoluted rulemaking process.

This abstract-sounding misalignment between the legal framework and the technical tools is not simply an academic worry, either. As the model development process and underlying algorithms become opaque to outsiders, well-financed stakeholders are better able to hijack the agency's agenda, dominating the analytical process and obscuring the many choices embedded within an algorithmic tool. Without a structured process for developing the algorithmic models and quality standards for underlying data sets, the agencies themselves may become rudderless and reactive to political and stakeholder pressures, allowing resulting algorithmic tools to be laden with errors and inadequate methods for evaluation. And, as noted, fishbowl transparency in these circumstances provides minimal protection since the high-tech tools require considerable expertise and resources. Yet even this fishbowl transparency is not always ensured, given the occasional proprietary claims surrounding some data sets, source-code, and other features of the algorithmic tools.[61]

1. REASON-GIVING

A central pillar guiding the work of agencies under the APA is the requirement that agencies provide reasons for their decisions and supporting analyses, and this reason-giving has been vigorously enforced by the courts. Courts have made it plain that in trying to extract these reasons (and hold agencies to account) they are endeavoring to motivate the agency not only to conduct a rigorous analysis to justify its work, but to do so ex ante—before the decision is made. As one of the early cases reminded agencies—reason-giving ideally should not amount to post hoc rationalization.[62]

But mapping this legal reason-giving aspiration against the ideal use of algorithmic tools exposes several sources of misalignment between the law and the real world of algorithmic modeling and other AI-based tools.

As an initial matter, reason-giving is generally oriented towards explaining the output or final "proposal" offered by the agency (the "what"), rather

than explaining and defending the rigor of the agency's analytical process underlying the rule (the "how").[63] Accordingly, the agency generally is not expected to explain *how* it develops a model—such as the process it used to develop the framing, choice of data, algorithms, and model selection. For example, in justifying why its emission limit for chemical tank farms is scientifically and legally sound, the Environmental Protection Agency will detail its comprehensive research of the relevant technological studies and its analytical decisions of how that evidence supports its legal standard. But the actual decision process itself—for example, the steps in framing its analytical work, which stakeholders the agency consulted at each step, the role of political officials at each step, and the role and selection of independent peer reviewers—is rarely detailed.[64] While occasionally the integrity of the agency's decision process itself (the how) is given a "hard look" or even cataloged as positive attributes of the agency's decision,[65] courts reverse agency rules only based on APA-triggered requirements, like providing a rulemaking record or offering an opportunity to comment.[66]

If reason-giving stopped short of demanding a type of stepwise explication of the underlying analysis, particularly for complex AI model development processes, the consequences might not be terribly problematic. Agency staff might still follow best practices voluntarily, engaging experts and the broader public at each step of the process.[67] But, as detailed below, administrative law actively discourages agencies from incorporating this structured process in its decisions for a number of overlapping reasons; in fact, agencies are sometimes legally better off avoiding this careful explication work.

How could this be?

First, under administrative law, agencies that do develop and follow a structured decision process—what Kevin Stack and Gillian Metzger call "internal law"—can find themselves more vulnerable to APA challenges rather than less.[68] Among the reasons for this counterintuitive result is that these added procedures become additional attachment points that can be used against (but not explicitly in favor of) the agency's decision. As Stack and Metzger document, "[t]hrough several different doctrinal routes, courts have gradually occluded the APA's openings for internal administrative law and transformed internal measures into externally enforced constraints."[69] Some of these internal law structures may be challenged outright as "notice

and comment rules" that require their own APA-compliance as legislative rules.[70] Other internal procedures may be characterized as legally binding, with agency deviations leading to actionable "noncompliance."[71]

Second, the enforcement of the reason-giving requirement in practice rests on an "on demand" approach that operates ex post, in reaction to errors or gaps, rather than ex ante to encourage excellent deliberative processes for the development and use of these tools.[72] One cannot claim the agency violated the APA by publishing a proposed rule without sufficient explanation (or reasons); that claim can only occur after the final rule is published and the agency still has not explained the challenged facts or choices after being prodded with specific comments to do so. If the agency's decision process itself lacked structure in unreasonable ways or was difficult to re-create, this is outside the range of challenges allowed under the APA.[73] Reason-giving, in terms of the actual, enforceable requirements, is hence "episodic" rather than comprehensive, predictable, and far-reaching.[74]

Third, only comments raised during notice and comment can be litigated.[75] As elaborated next, if the comments are one-sided, so will be the oversight of the agency's reason-giving. Regulatory participants that work with the agencies to develop algorithmic tools—or offer private tools to save the agency the trouble—may be the only commenters engaged in a rule.

Finally, "[f]orego[ing] clarity and publicity in order to avoid having to defend their policies and decisions, or to preserve maximal room to maneuver in the future" offers a number of discretionary advantages to agency decisions that might tip the balance towards taking full advantage of these various reason-giving loopholes.[76] Political operatives, as well as agency managers, will appreciate that less structure means more elbow room to accommodate squeaky wheels, whoever they are. Moreover, as long as the focus is on substance and not process, the possibility for ends-oriented (post-hoc) rationalizations become more likely. If the algorithmic tool is producing unwelcome results, for example, political officials could intervene in ways that alter the way the tool is implemented, developed, or how the tool is explained. There are generally no barriers to protect against this rigging of the technical analysis when the White House has an interest.[77]

2. DEMOCRATIC DELIBERATIONS

Even if the development and use of algorithmic tools occurred without best-practice types of explications and structures, the APA seems well-suited, at least on paper, to bring skepticism to bear from diverse parties. Not only must the agency publish its proposal and accompanying analysis (e.g., choice of algorithmic tools, data sets) for review by all members of the public, including experts, but it must "consider" these comments under threat of judicial review. Thus, in theory, rigorous public engagement should catch significant bias, manipulations of algorithmic processes, incompetence, poor framing choices, and other agency indiscretions.

But, in practice, the APA-mandated process falls quite short of ensuring this diverse oversight.[78] Instead of conditioning rules on truly "democratic deliberation," administrative process requires agencies to only provide an *opportunity* for stakeholders to voluntarily review and comment on an agency's analysis and proposal. There are no subsidies to ensure this participation represents all affected interests. The agency is also not required to reach out and solicit the input of unrepresented groups, much less certify that participation has been comprehensive and diverse.

Hence, when the costs to participate grow too high, some subsets of participants—particularly those who are thinly financed, like many public interest groups—drop out of the process.[79] As a result, the "democratic" oversight of agencies in practice can operate on a type of pay-to-play system. Those who can afford to engage become active and influential participants; those who lack the expertise, time, or resources are unrepresented in notice and comment, and their interests can be ignored. Even during judicial review, courts do not consider whether public participation was diverse or generally representative; the court focuses instead on the particularized challenges lodged by the litigants before it.

Abstract worries that the resulting notice and comment processes might be afflicted with significant democratic deficits is borne out in empirical research. A number of empirical studies of rulemakings identify a "bias towards business" that runs throughout the rulemaking process.[80] Industry groups tend to dominate and sometimes monopolize both the development of the rule proposal, and the notice and comment process. See Figure 3.

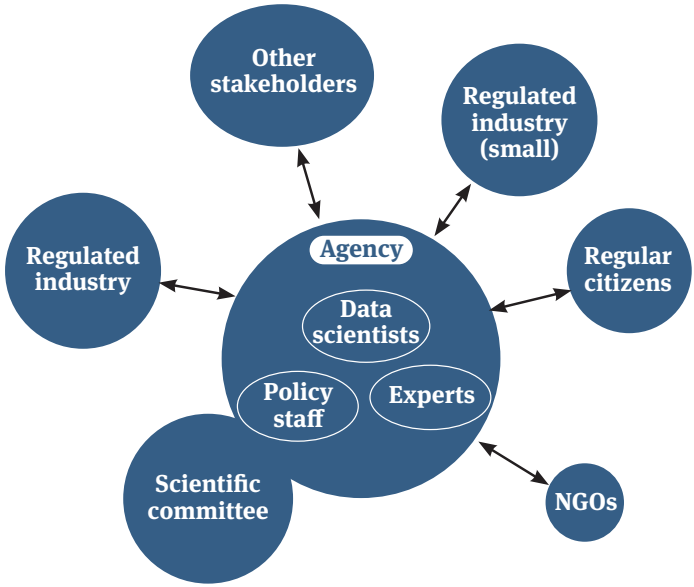
These risks of a democratic deficit are particularly acute for agency

decisions supported by algorithmic tools.[81] Theorists have long noted how the "costs" and "complexity" of the regulatory issues directly affect the types of participants who can realistically engage in regulatory deliberations.[82] As the costs or complexity of the decision or supporting evidence rise, certain affected groups will find they can no longer afford to participate.[83] Without concerted educational programs and outreach, as noted in Section I, algorithmic tools will generally not be amenable to input or review by these nonexperts.

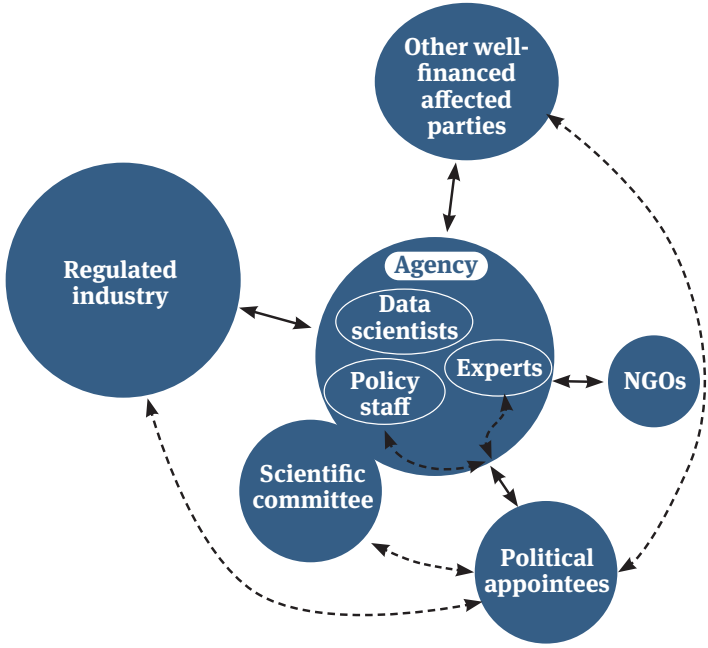
But the passive approach to participation can lead to even more dramatic biasing effects on the development of complex algorithmic tools. First, based on the courts' "logical outgrowth test," which they have read into the APA, agencies understand that if comments and deliberations reveal that material changes are needed to its initial proposal, the agency must re-start the notice and comment process over in order to ensure meaningful opportunities for participation.[84] From the agency's perspective, then, it is far better to propose a rule that is effectively complete at the initial "proposal" stage than risk being blind-sided during notice and comment.[85]

Shrewd agencies eager to get a rule through the legal process will find that they can limit these litigation risks if they reach out to stakeholders—particularly litigious, vocal stakeholders—early in the decision process in order to anticipate and address the major concerns in advance, to minimize the risks of judicial review.[86] (This technique may be equally useful for agency policies not subject to strong APA requirements since it helps placate some concerns by political appointees.) Administrative law does not restrict these types of pre-proposal deliberations or even require agencies (in most programs) to log these early-stage communications. "Negotiating" algorithmic choices, data choices, evaluative metrics, and other algorithmic features with a subset of stakeholders before a rule is proposed is thus fully in bounds of administrative law.[87]

Second and more perverse, if the agency's proposal is effectively incomprehensible in ways that alienate or unduly raise the costs to thinly financed stakeholders' ability to make sense of the proposal, the APA's unintended response is to reward (rather than sanction) the agency with fewer comments and litigation risks.[88] As just discussed, a number of affected groups can find themselves priced out of these rulemaking deliberations precisely because



(a) An idealized mapping of participation in rules.



(b) An actual mapping of participation in rules as discussed in this essay.

FIGURE 3: A SCHEMATIC COMPARING IDEAL VS. ACTUAL PATTERNS OF STAKEHOLDER PARTICIPATION IN ENVIRONMENTAL REGULATION. DASHED LINES INDICATE POSSIBLY NON-TRANSPARENT RELATIONSHIPS

the costs of making sense of the agency's convoluted algorithmic processes, choices, and analysis is simply too high. But in a pragmatic sense, their inability to participate translates to fewer legal vulnerabilities for the agency. In *Bending the Rules*, Rachel Potter documents how agencies sometimes do act strategically by drafting rule proposals and underlying analyses in ways that make them more inscrutable to outsiders.[89]

Finally, those groups that do participate in the regulatory deliberations have the opportunity to exert disproportionate influence over the agency's analysis and reasoning as a result of the law's requirement that the agency "consider" all submitted comments.[90] Precisely because there is no limit to the number and size of these comments, participants can (and sometimes do) take advantage and strategically inundate the agency with minutiae, ends-oriented critiques, and even prepare their own, alternative analyses, which, in the case of AI tools, could involve introducing their own preferred AI-based model choices, algorithm implementations, and training data sets, as well as the utilization of proprietary modeling packages.[91] Battles over each individual piece of code or each data set launched during notice and comment thus threaten to overwhelm and drown out more fundamental problems with the algorithmic tools that are in need of input and critique, including even the purpose of the algorithmic tool itself.[92] These "blunderbuss" techniques by strategic stakeholders have the added benefit of making the agency's analyses and algorithmic tools still more convoluted and inaccessible to outsiders.

3. EXPERT REVIEW

Expert peer review provides another partial corrective, in theory, to guard against some inscrutable and biased algorithmic tools used to inform regulation. While expert peer reviewers have no business making the many public choices embedded in algorithmic tools—like how and how much to spend protecting the public from risks—these experts can be (and have been) quite valuable in extracting the numerous junctures occurring within a model that necessitate public guidance and input, which might have otherwise escaped public scrutiny. Expert review is also helpful in isolating occasions in which the model's algorithms, data sets, or design may lack sufficient documentation or drift from consensual scientific standards. As a result of

these contributions, expert peer review has become a central ingredient for most of the agencies' technical decision-making and could play a helpful role in overseeing some algorithmic tools as well (although identifying the relevant types of experts may be more difficult).

But, once again, the law governing a slice of this external peer review—formal expert advisory panels are governed by the Federal Advisory Committee Act (FACA)—provides room for political manipulation of technical decisions, including algorithmic tools.[93] This occurs because the selection of scientists for these peer review panels, as well as many of the initial decisions to create (or continue) external science advisory panels and to frame their charges, rests ultimately with political officials.[94] And, although in general it appears this political control has not been abused, that is not always the case.[95] The Trump administration even terminated or limited the role of expert panels when their views were expected to expose flaws in politically crafted technical analyses supporting a preordained regulatory decision.[96]

4. UNDERUTILIZATION

Even if the development and use of algorithmic tools may sometimes drift away from best practices in worrisome ways, the resulting algorithmic tools might still offer vital analyses that, while not perfect, are better than nothing.[97] But again, administrative law produces disappointments since, in practice, administrative law is structured to encourage agency underutilization, rather than overutilization, of these tools. More specifically, agencies can be sued for promulgating "arbitrary" rules, but if the authorizing statute does not specify deadlines, the fact that the agency does nothing is generally unreviewable.[98] The resulting risks of underutilization seem particularly acute for algorithmic tools because they require more choices, experimentation, and risk that are vulnerable to challenge.[99] The "reason-giving" requirements of administrative law can be used strategically against algorithmic tools to capitalize on ways that they are always "wrong" and hence contestable, particularly when the referee is a nonexpert court.[100] In an administrative law system that does not "require" the use of algorithmic tools on behalf of the public, but subjects the agencies that do attempt to develop these tools to unlimited forms of critical attack by well-financed groups that stand to lose from the innovation, the rational course for agencies is clear.

As a result of these reinforcing legal costs, the "sandbox" of experimentation, particularly advocated for with AI tools, is more likely to become a graveyard of big data projects that could not survive the firestorm of critical commentary.[101] Algorithmic processes that have the capacity to be of great value in advancing public goals—e.g., identifying environmental hazards and potential violators—may be especially discouraged, at least when their use runs against the interests of well-funded stakeholders. For example, at least four separate regulatory programs in the environmental and public health arena require "adverse events" reports occurring post-market on drugs, chemicals, pesticides, etc., from manufacturers.[102] Yet despite evidence that these adverse events reports are historically insufficiently analyzed by staff due to a lack of resources, none of the relevant agencies—except an early pilot effort at FDA—are preliminarily exploring the use of tools like AI to assist in screening them.[103] Indeed, despite the many touted applications of AI to environmental and public health decision-making, the ACUS study identified only one EPA use of AI (out 157 government-wide projects)—for the CompTox Chemicals Dashboard.[104]

5. MAYBE IT WILL ALL WORK OUT OK ANYWAY?

This grim account of perverted reason-giving and a democratic deficit arising from the design of administrative law may not always come to pass. The agency is still filled with experts who are likely to try to persevere regardless of the legal impediments.[105] On the other hand, understaffing, political interventions, and attrition of top officials reveal that civil staff is not an automatic, magic ingredient to the successful, appropriate, and fair use of algorithmic tools.[106]

Additionally, in some settings there will be sufficient variation among participating stakeholders to lead to much more robust deliberations, including over algorithmic tools. Once the participants are diverse and vigorous, the opacity surrounding algorithmic tools is more likely to yield to the light (and the courts), as the warring participants debate specific features.[107] But even if this constructive adversarial scrutiny occurs in some rules, it might not change the agencies' decision process across all rules (including those involving one-sided participation).

Finally, it is possible that many valuable uses of algorithmic tools by agencies will occur outside of the rulemaking context. However, this

possibility offers less, rather than more reason for comfort since it means that the deliberative, reason-giving requirements of administrative process might be circumvented altogether. Well-financed, organized stakeholders may find it even easier to exert control over the use of algorithmic tools in settings where the agency is not required to solicit public comment.

III. REFORM AND RECOMMENDATIONS

OUR ANALYSIS IDENTIFIES a number of ways that administrative law appears ill-prepared to encourage the accountable use of algorithmic tools, at least in environmental and public health regulation. These findings, which are consistent with those of other scholars, reinforce the need for legal attention.[108]

In this section, we suggest a suite of reforms designed to adjust misaligned legal requirements and provide support for agency excellence in the use of algorithmic tools.

A. Top down: Best practices encouraged

The preliminary best practices outlined in Section I draw on early and emerging consensus on how algorithmic tools *should* be developed and used to inform policy. Yet, as discussed, current legal interpretations of the APA are not only silent on, but inadvertently work to tacitly discourage agencies from instituting these best practices in the development and use of algorithmic tools.

At first blush, the most straightforward reform for this problem would simply be to require agencies to follow these best practices in their development and use of algorithmic tools. However, there are numerous reasons why imposing more legal requirements on agency practices, particularly in this dynamic and technical area, are ill-advised.[109] Most significant is the risk that legal requirements will be co-opted by well-financed parties as added attachment points for litigation that distract and wear down the agencies. If agencies are already discouraged from innovating with algorithmic tools in part due to concern about added legal vulnerabilities, creating more mandatory requirements will not help the situation. Mandating best practices

also raises very real dangers of an ill-fitting, one-size-fits-all approach to this fast-moving area of data science that might inhibit creativity and innovation. Or, still worse, the legal requirements themselves might be crafted by politicians with an aim to bias the process in unaccountable ways from the start.[110]

1. LEGAL REWARDS FOR BEST PRACTICES

Instead, we advocate legal rewards for agencies using best practices in a meaningful and rigorous way that count only in favor of an agency's analysis and not against it. One possible reform along these lines is an amendment to the APA that would direct courts to be highly deferential—perhaps a "manifestly erroneous" standard—when the agency convincingly demonstrates it developed an algorithmic tool in a way that engaged publics and experts meaningfully on the key choices.[111] (Statutory interpretation questions would obviously remain unaffected.)

The agency's showing of an "accountable decision process" would thus serve as an affirmative defense to arbitrary and capricious challenges that, when successful, would lead to a much more deferential review. As a result of this defense, if an agency takes the initiative to create and document an excellent, structured decision process to develop its algorithmic tool, it would be insulated from most litigation risks associated with its technical analyses. Accountable decision processes thus become a benefit, rather than a cost to a rational agency. Moreover, with all key choices out in the open, other forms of oversight—particularly by the political branches—will be able to engage to keep the agency's choices rigorous and accountable.

2. INTERNAL LAW: CREATION OF A CENTRALIZED, FEDERAL EXPERT OFFICE

A second reform focuses on shoring up the agencies' "internal law" governing algorithmic tools. To that end, we recommend the creation of a pragmatically oriented federal office or institute that, among its functions, provides assistance to agencies in ensuring inclusive, democratic engagement in the algorithmic tools used for policy. Some members of Congress have already recommended the creation of a centralized, federal expert office to assist and encourage agencies' use of algorithmic tools, but these congressional proposals are generally silent on the need for the office to also assist agencies with

the very challenging task of ensuring meaningful public participation.[112]

In terms of general function, the new office would encourage greater agency innovation in the inclusive development of algorithmic tools in the policymaking process by offering one-to-one expert staff assistance and developing a wealth of technical resources to catalyze interagency partnerships. The office might even provide an internal, expert "certification" process that reviews particularly important agency uses of algorithmic tools consistent with its own guidelines—particularly regarding engaging meaningful public participation—to provide the tools with more insulation against unfair stakeholder attack and political pressure. The office would be an important resource for providing valuable input to the National Artificial Intelligence Initiative (NAII) and the Center for Strategic Foresight of the U.S. Government Accountability Office (GAO), among others.

3. PROVENANCE AS A CROSS-CUTTING ISSUE

Decisions are made at each stage of the algorithmic tool development process, including which stakeholders should be included, whether the agency should adopt AI for a task, the reasons for the adoption of specific algorithms, the choice of data sets, etc. This chain of decisions is the soul of the final algorithmic tool.[113]

Agencies should thus be expected to provide a provenance statement that documents all such decisions taken throughout the development and use of algorithmic tools, as well as identifies all the contributors and reviewers who contributed at each step.[114] Building on the scientific journal's use of authorship and conflicts of interest disclosures, for example, each contributor and reviewer could also be identified with at least this information, along with a link to his or her specific contribution in the larger project. Provenance statements could be required through an executive order, adopted informally by agencies, or might even be required through legislation.

B. Ensuring diverse participation by experts and publics

Ensuring that all affected stakeholders and diverse experts are adequately consulted in the development of algorithmic tools will require additional adjustments to the administrative process. We begin with the most difficult

challenge first—engaging meaningful public participation.

1. PUBLIC PARTICIPATION IN AGENCIES' USE OF ALGORITHMIC TOOLS

Even with structured, accessible decision processes and the disclosure of provenance, a passive approach to public engagement may still lead to underrepresentation of thinly financed groups in the development and oversight of algorithmic tools to inform policy decisions. One of the goals of best practices is to infuse balanced and diverse participation by the public throughout the process. To that end, we offer two parallel tracks of recommendations to ensure more complete, inclusive public engagement.

First, the agency should map—ex ante—where the most vital roles for public engagement lie in the development of a given algorithmic tool, and then use that map to actively solicit meaningful and diverse public deliberations on key issues. In at least one EPA program (setting of ambient air quality standards), the EPA provides a preliminary template for how this could be done. In that program, the agency conducts an initial "scoping" session in an open meeting with policymakers and any interested publics. (Ideally, the agency would actively solicit and even subsidize participation to ensure that all relevant views are present). This scoping session results in a written, elaborate "planning" document that frames the questions and public preferences to guide subsequent technical analyses and models.[115] The agency concludes its use of models with a final, scientifically peer-reviewed report that has as its sole purpose "translating" the scientific work into terms that lay persons, including policymakers and attentive publics, would understand.[116]

In identifying targeted opportunities for public input and engagement, the agency should also design the deliberations to be meaningful. Rather than a pro forma fishbowl-transparency approach to participation, for example, the agency should engage the full range of stakeholders, even if some outreach is needed to ensure the broader publics' views are adequately represented in the deliberations. Indeed, this showing of inclusive, meaningful deliberations could then be part of the affirmative defense for the "excellent decision processes" just discussed to legally reward (and encourage) agencies for more inclusive regulatory deliberations.

Second, it may be necessary to provide subsidies to support participation by thinly financed affected stakeholders and/or institutionalize publicly appointed advocates to engage on their behalf. Given the institutional inertia that appears to favor the agencies' underutilization of algorithmic tools for the public interest, these subsidies and supports for the broader public should not only target algorithmic tools undergoing development but should reach early into the agencies' decision process to ensure that the tools are not being sidelined or overlooked when they could be used to solve intractable data shortages and other problems.

2. EXTERNAL EXPERT OVERSIGHT

There are ongoing disagreements about the normative desirability of political control over FACA-based science advisory boards, as well as what shape reform might take.[117] We simply flag the potential biasing effects that the current FACA approach could have in agencies' development of rigorous and accountable models. Given the narrow scope of our study, however, we must save this larger reform project for other analysts.[118]

C. Bottom-Up: Privately produced algorithmic tools used by agencies

Agencies will not be the only parties preparing algorithmic tools for use in policy. Regulated parties, trade associations, think tanks, educational institutions, and some public groups will also innovate in developing their tools, data sets, and corresponding models designed to inform policy.[119]

1. MAINSTREAM APPROACHES

In using privately developed algorithmic tools for regulation,[120] the private sponsor/author should be required to demonstrate full compliance with a rigid set of best practices. As the use and complexity of AI-generated models increase, the burden of evaluating privately produced algorithmic tools will quickly become far too high to place on agencies and other, skeptical stakeholders. By placing the burden on private sponsors/authors to document algorithmic tools' integrity as a condition to their use in regulation, legal responsibility is placed on the party best able to do this work.

2. OPENING THE DOOR TO PUBLIC STAKEHOLDER PARTNERSHIPS IN DEVELOPING ALGORITHMIC TOOLS

In some settings, agencies may utilize proprietary algorithmic tools provided by private organizations, in spite of the fact that these tools cannot be scrutinized by all participants due to intellectual property protections. When this occurs, agencies should be required to subsidize public stakeholders to develop their own, open source algorithmic tools that could be used in place of the proprietary tools. Stakeholders could involve partnerships of civil society actors, educational institutions, and regulated entities. This approach should help counteract the ability of private technology companies to monopolize policy deliberations using proprietary tools. Facilitating these alternative, public models also tracks best practices by incorporating added hypotheses and alternatives into the decision process.

D. General recommendation

The adoption of algorithmic tools throughout the vast U.S. bureaucracy will not occur overnight. Agencies are confronting a number of new challenges, including building in-house expertise, locating or creating reliable data sets, and adjusting their decision processes.[121] Moreover, since many agencies are considerably ahead of others and their algorithmic needs are different, there is no one-size-fits-all approach.

Therefore, going forward, it is important for the U.S. to institute two tracks for assessing the agencies' use of algorithmic tools. The first track would consist of an ongoing, ACUS-styled study of agencies' use of these algorithmic tools over time—perhaps focusing on algorithmic tools used by a few representative agencies as set against the best practices. The second track would assess all agencies' readiness to absorb AI and recommend how to better prepare them to use these new tools in effective and accountable ways.[122]

Both assessments will likely flag new, unexpected challenges and problems occurring within agencies, as well as showcase ways that agencies are rising to the challenges by developing innovative ways to ensure the accountability and public benefits of their algorithmic tools. Only through this "real-time" assessment can we be sure we are making the most of these new technologies, while simultaneously protecting against the very real risk

that some of these tools might be used in ways that undermine the vital goal of accountable agency decision-making.

IV. CONCLUSION

Decades of experience using mechanistic models to support environmental and public health regulations expose numerous cracks in the administrative law system, with respect to encouraging the accountable use of complex, technical tools. As the applications of AI/ML grow, these cracks in the system could become fissures. Fortunately, emerging consensual best practices drawn from the literature help identify means for significantly improving the accountable development of algorithmic tools for regulation. We offer preliminary recommendations for how we might better align legal rules with these best practices, in order to encourage and support the agencies' uses of these new technologies to advance the public interest in the future. Are we asking too much? Given the risks to basic democratic values from the unwise or inadequate use of these new, important tools in regulation, the need for targeted legal reforms seems well worth the investment.

NOTES

1 *E.g.*, Environmental Law Institute, *Big Data and Environmental Protection: An Initial Survey of Public and Private Initiatives* (2014), https://www.eli.org/sites/default/files/eli-pubs/big-data-and-environmental-protection.pdf [https://perma.cc/RW9M-R9TR]; Claire Lajaunie, et al., *Big Data Enters Environmental Law*, 8 Transnat'l Env't L. 523 (Nov. 2019), https://doi.org/10.1017/S2047102519000335 [https://perma.cc/56J7-9REQ]; J. Wambaugh, Nat'l Ctr. for Computational Toxicology, Off. Rsch. and Dev., et al., Predicting Exposure Pathways with ML at the International Society of Exposure Science Annual Meeting (Oct. 15-19, 2017), https://cfpub.epa.gov/si/si_public_file_download.cfm?p_download_id=536325&Lab=NCCT [https://perma.cc/3HEU-K5QS].

2 *E.g.*, Antony Williams et al., *The CompTox Chemistry Dashboard: a community data resource for environmental chemistry.* 9 J Cheminform 61 (2017).

3 *E.g.*, Robert L. Glicksman et al., *Technological Innovation, Data Analytics, and Environmental Enforcement*, 44 Ecology L. Q. 41, 64-86 (2017).

4 *E.g.*, Zachary Brennan, *FDA Seeks Help Using Algorithms to Detect Adverse Event Anomalies, Regulatory Affairs Professionals Society*, Regul. Aff.'s Pro.'s Soc'y (Jan. 23, 2020), https://www.raps.org/news-and-articles/news-articles/2020/1/fda-seeks-help-using-algorithms-to-detect-adverse.

5 *See generally* Wendy Wagner, Elizabeth Fisher & Pasky Pascual, *Misunderstanding Models in Environmental and Public Health Regulation*, 18 N.Y.U. Envtl. L.J. 293 (2010); Pasky Pascual, Wendy Wagner & Elizabeth Fisher, *Making Method Visible: Improving the Quality of Science-Based Regulation*, 2 Mich. J. Envtl. & Admin. L. 429, 470–71 (2013).

6 We thus focus on a different, complementary set of agency decisions as compared with Engstrom and Ho's study of AI in the regulatory enforcement and adjudication context. *See generally* David Freeman Engstrom & Daniel E. Ho, Algorithmic Accountability in the Administrative State, 37 Yale J. on Reg. (2020), https://digitalcommons.law.yale.edu/yjreg/vol37/iss3/1/ [https://perma.cc/RYK8-T7GC].

7 *E.g.*, Finale Doshi-Velez et al., Accountability of AI Under the Law (undated) (working draft), (https://arxiv.org/ftp/arxiv/papers/1711/1711.01134.pdf [https://perma.cc/RPV2-XTR3]).

8 *E.g.*, Mariano-Florentino Cuéllar, *Cyberdelegation and the Administrative State, in* Administrative Law from the Inside Out: Essays on Themes in the Work of Jerry L. Mashaw 134 (Nicholas R. Parrilo ed., 2017).

9 *E.g.*, Martin Murillo & S. Naidu, Discrete-time optimal control systems with state constraints at the AIAA Guidance, Navigation, and Control Conference and Exhibit (2002), https://arc.aiaa.org/doi/10.2514/6.2002-4665 [https://perma.cc/GD2Z-Q2PE].

10 *E.g.*, John Nay & Katherine J. Strandburg, *Generalizability: Machine Learning and Humans-in-the-loop, in* Research Handbook on Big Data Law 2 (Roland Vogl, ed., forthcoming 2021).

11 Katherine J. Strandburg, *Rulemaking and Inscrutable Automated Decision Tools*, 119 Colum. L. Rev. 1851, 1862 (2019).

12 Frank Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information (2015).

13 *E.g.*, Joshua A. Kroll et al. *Accountable Algorithms*, 165 U. Pa. L. Rev. 638 (2017).

14 Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 Admin. L. Rev. 1 (2019).

15 *E.g.*, Id.; Kroll et al., *supra* note 13.

16 *See generally* Sheila Jasanoff, *Serviceable Truths: Science for Action in Law and Policy*, 93 Tex. L. Rev. 1723 (2015).

17 Kroll et al., *supra* note 13, at 636-37, 702.

18 Coglianese & Lehr, *supra* note 14, at 19.

19 *E.g.*, Pasquale, *supra* note 12, at 16; Coglianese & Lehr, *supra* note 14, at 19.

20 *E.g.*, Kroll et al., *supra* note 13

21 Naomi Oreskes, Why Trust Science? (2019).

22 Kroll et al., *supra* note 13; Katherine J. Strand-

burg, *Rulemaking and Inscrutable Automated Decision Tools*, 119 Colum. L. Rev. 1851, 1862 (2019).

23 *See generally* Wagner, Fisher & Pascual, *supra* note 5.

24 National Research Council, National Academies of Sciences, *Models in Environmental Regulatory Decisionmaking* (2007) [hereinafter NAS Models].

25 David F. Engstrom & Daniel E. Ho, *Artificially Intelligent Government: A Review and Agenda*, *in* Research Handbook on Big Data Law 14-15 (Roland Vogl ed., forthcoming 2021), (https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3551549 [https://perma.cc/6LEN-55WJ]); Kroll et al., *supra* note 13, at 637, 640; *cf.* Strandburg, *supra* note 22, at 1872-73, 1882-83; NAS Models, *supra* note 24, at 4.

26 Kroll et al,, *supra* note 13, at 640.

27 *Id.* at 637-38.

28 Engstrom & Ho, *supra* note 25, at 6-7 (citing Finale Doshi-Velez & Mason Kortz, *Accountability of AI Under the Law: The Role of Explanation* (Berkman Klein Ctr., Working Paper, 2017)).

29 David F. Engstrom, et al., ACUS Study: Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies at 75 (2020) [hereinafter ACUS AI Study].

30 Strandburg, *supra* note 22, at 1855.

31 NAS Models, *supra* note 24, at 88-98; Nay & Strandburg, *supra* note 10, at p.13, 18; Kroll et al., *supra* note 13, at 656, 701, 702-703; Strandburg, *supra* note 22, at 1880.

32 *E.g.*, Jasanoff, *supra* note 16.

33 One possible representation of a policy need is a Requirements Specification document, See: 830-1984 - IEEE Guide for Software Requirements Specifications, https://ieeexplore.ieee.org/document/278253 [https://perma.cc/WQ9G-6WFL].

34 NAS Models Study, *supra* note 24, at 98-99.

35 Engstrom & Ho, *supra* note 25, at 6.

36 Even in the extreme case that policy tools are provided in "as a service" framework, stakeholders should be able to utilize them for assessment, simulation, benchmarking, and other purposes.

37 Kroll et al., *supra* note 13, at 672; Coglianese & Lehr, *supra* note 14, at 34-36.

38 Kroll et al., supra note 13, at 680-82.

39 *E.g.*, Nay and Strandburg, *supra* note 10, at 1.

40 NAS Models Study, *supra* note 24, at 4.

41 Kroll et al., *supra* note 13, at 682.

42 NAS Models Study, *supra* note 24, at ch.4; *see also* Kroll et al., *supra* note 13, at 672-674, 677.

43 The ACUS study reports that 61% of government AI applications present insufficient publicly available technical documentation to even determine what methods are being deployed. ACUS AI Study, *supra* note 29, at 19.

44 Coglianese & Lehr, *supra* note 14.

45 See also Wagner, Fisher & Pascual, *supra* note 5.

46 Engstrom & Ho, *supra* note 6, at 7; *see also Id.* at 28, 38, 44; *see also* Kroll et al., *supra* note 13, at 636.

47 Strandburg, *supra* note 22, at 1852; Engstrom & Ho, *supra* note 6, at 5.

48 Strandburg, *supra* note 22, at 1880; *see also* NRC, Review of the EPA's Draft IRIS Risk Assessment of Formaldehyde, ch. 7 (2011).

49 *E.g.*, Jerry Mashaw, Reasoned Administration and Democratic Legitimacy (2018); Mark Seidenfeld, *A Civic Republican Justification for the Bureaucratic State*, 105 Harv. L. Rev. 1511, 1514 (1992).

50 5 U.S.C. §§ 551 et seq.

51 *Id.* at § 706(2)(a); *see also* Motor Vehicle Mfrs. Assn'n v. State Farm Mut. Auto. Ins. Co., 463 U.S. 29. 43 (1983).

52 Edward Rubin, *It's Time to Make the Administrative Procedure Act Administrative*, 89 Cornell L. Rev. 95, 101 (2003).

53 Michael Sant'Ambrogio & Glen Staszewski, *Democratizing Rule Development*, 98 Wash. Univ. L. Rev 793 (2021), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3615896 [https://perma.cc/H9UA-C2CA].

54 Emily Hammond Meazell, *Super Deference, the Science Obsession, and Judicial Review of Agency Science*, 109 Mich. L. Rev. 733, 738, 778-79 (2011).

55 *Id.*

56 Louis J. Virelli III, *Deconstructing Arbitrary and*

Capricious Review, 92 N.C. L. Rev. 721 (2014).

57 Sant'Ambrogio & Staszewski, *supra* note 53.

58 Wendy E. Wagner, *A Place for Agency Expertise: Reconciling Agency Expertise with Presidential Power*, 115 Colum. L. Rev. 2019, 2032–45 (2015).

59 Nina A. Mendelson, *Disclosing "Political" Oversight of Agency Decision Making*, 108 Mich. L. Rev. 1127, 1157 (2010).

60 *E.g.*, Brad Plumer & Coral Davenport, *Science Under Attack: How Trump Is Sidelining Researchers and their Work*, N.Y. Times, Dec. 29, 2019 at A1.

61 Hannah Bloch-Wehba, *Access to Algorithms*, 88 Fordham L. Rev. 1265, 1273-89 (2020).

62 Citizens to Preserve Overton Park, Inc. v. Volpe, 401 U.S. 402, 419-421 (1971).

63 *Cf.* Vermont Yankee Nuclear Power Corp. v. Natural Resources Defense Council, 435 U.S. 519 (1978).

64 *See generally* Wendy Wagner, Incomprehensible! 162-68 (2019)

65 *E.g*, *Id*. at 45-47.

66 Virelli, *supra* note 56, at 738.

67 *Cf.* Engstrom & Ho, *supra* note 25, at 14.

68 Gillian E. Metgzer & Kevin M. Stack, *Internal Administrative Law*, 115 Mich. L. Rev. 1239 (2017).

69 *Id*. at 1245.

70 *Id*. at 1281.

71 *Id*. at 1288.

72 *Id*. at 1264.

73 *Id*.

74 *Id*. It may also be relatively infrequent, Wendy Wagner, *Revisiting the Impact of Judicial Review on Agency Rulemaking: An Empirical Investigation*, 53 Wm. & Mary L. Rev. 1717-95 (2012), and quite deferential in judging these reasons. *E.g.*, Adrian Vermeule, Law's Abnegation 34 (2016).

75 NRDC v. Thomas, 805 F.2d 410 (D.C. Cir. 1986).

76 Metzger & Stack, *supra* note 68, at 1303.

77 Thomas McGarity & Wendy Wagner, *Deregulation Using Stealth Science Strategies*, 68 Duke L. J. 1719, 1724–34 (2019).

78 This argument is laid out more fully in Wagner, *supra* note 64, at ch. 6 (2019).

79 Cynthia R. Farina, Mary J. Newhart & Cheryl Blake, *The Problem with Words: Plain Language and Public Participation in Rulemaking*, 83 Geo. Wash. L. Rev. 1358, 1365 (2015). We necessarily must bracket issues concerning how the public is represented and instead assume that there will likely be NGO experts who purport to represent the general public. *E.g.*, Strandburg, *supra* note 22, at 1875; *but see* Miriam Seifter, *Second-Order Participation in Administrative Law*, 63 UCLA L. Rev. 1300 (2016).

80 *E.g.*, Jason W. Yackee & Susan W. Yackee, *A Bias toward Business? Assessing Interest Group Influence on the Bureaucracy*, 68 J. Pol. 128 (2006); *see* literature cited in Wagner, *supra* note 64, at 190-91.

81 *E.g.*, Engstrom & Ho, *supra* note 6, at 37.

82 *E.g.*, William T. Gormley, Jr., *Regulatory Issue Networks in a Federal System*, 18 Polity 595, 607 (1986).

83 *Cf.* Rachel Potter, Bending the Rules 70 (2019).

84 Shell Oil Co. v. EPA, 950 F.2d 741, 757–63 (D.C. Cir. 1991).

85 William F. West, Formal Procedures, *Informal Processes, Accountability and Responsiveness in Bureaucratic Policy Making: An Institutional Policy Analysis*, 64 Pub. Admin. Rev. 66 (2004), https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6210.2004.00347.x [https://perma.cc/BN58-UV8B].

86 *Id*.

87 NRDC v. EPA, 595 F.Supp. 1255, 1259 (S.D.N.Y. 1984).

88 Wagner, *supra* note 64.

89 Potter, *supra* note 82, at 70-74.

90 Wagner, *supra* note 64, at 186-87.

91 *Id*.

92 Thomas McGarity, *Defending Clean Science from Dirty Attacks by Special Interests*, *in* Rescuing Science from Politics (Wendy Wagner and Rena Steinzor eds., 2006).

93 McGarity & Wagner, *supra* note 76, at 1756-57.

94 Robert Steinbrook, *Science, Politics, and Federal Advisory Committees*, 350 NEJM 1454, 1457 (2004).

95 *Id.*

96 McGarity & Wagner, *supra* note 76, at 1756-66.

97 *E.g.*, Strandburg, *supra* note 22, at 1881.

98 *E.g.*, Sidney A. Shapiro, *Rulemaking Inaction and the Failure of Administrative Law*, 68 DUKE L. J. 1805 (2019).

99 *E.g.*, ACUS AI Study, *supra* note 29, at 89.

100 *E.g.*, Engstrom & Ho, *supra* note 6, at 34.

101 ACUS AI Study, *supra* note 29, at 89.

102 Adverse effects are required to be submitted to EPA under FIFRA and TSCA; to the Consumer Product Safety Commission; to the Occupational Safety and Health Administration; and to the National Highway Traffic Safety Administration, as well as to the Food and Drug Administration.

103 *E.g.*, Zachary Brennan, *FDA Seeks Help Using Algorithms to Detect Adverse Event Anomalies*, REGUL. AFF.'S PRO.'S SOC'Y (Jan. 23, 2020), https://www.raps.org/news-and-articles/news-articles/2020/1/fda-seeks-help-using-algorithms-to-detect-adverse. *See also* ACUS AI Study, *supra* note 29, at 55-56, 90 (noting cross-agency challenges in adverse reports analyses and implying that even the FDA is dragging its feet by not requiring structured data by private parties).

104 *Id.* at Appendix, Case #106, https://law.stanford.edu/education/only-at-sls/law-policy-lab/practicums-2018-2019/administering-by-algorithm-artificial-intelligence-in-the-regulatory-state/acus-report-for-administering-by-algorithm-artificial-intelligence-in-the-regulatory-state/ [https://perma.cc/XH28-T64Q].

105 JOHN BREHM & SCOTT GATES, WORKING, SHIRKING, AND SABOTAGE: BUREAUCRATIC RESPONSE TO A DEMOCRATIC PUBLIC 79-81 (1999); MARISSA MARTINO GOLDEN, WHAT MOTIVATES BUREAUCRATS? POLITICS AND ADMINISTRATION DURING THE REAGAN YEARS 106–07 (2000).

106 The ACUS Study also identified a lack of technical sophistication in 88% of the AI applications. ACUS AI Study, *supra* note 29, at 7.

107 *E.g.*, Sidney Shapiro et al., *The Enlightenment of Administrative Law: Looking Inside the Agency for Legitimacy*, 47 WAKE FOREST L. REV. 463, 491–502 (2012).

108 *See also* Wagner, Fisher & Pascual, *supra* note 5. Additionally, although we examine a different portion of the universe of algorithmic tools, our conclusions are almost identical to those of Engstrom and Ho. *E.g.*, Engstrom & Ho, *supra* note 6, at 21; *see also* Kroll et al., *supra* note 13, at 636.

109 *See generally* Nicholas Bagley, *The Procedure Fetish*, 118 MICH. L. REV. 345 (2019).

110 *E.g.*, Wendy Wagner et al., *Whose Science? A New Era in Regulatory Science Wars*, 362 SCI. 636-39 (2018)

111 *See also* WAGNER, *supra* note 64, at 195-99.

112 *See* summary in Lee Tiedrich et al., *Recent Developments in Artificial Intelligence Law and Policy*, REGUL. REV. (Apr. 25, 2020), https://www.theregreview.org/2020/04/25/tiedrich-altvater-yoon-recent-developments-artificial-intelligence-law-policy/ [https://perma.cc/TGH8-UDQY]; *see also* Engstrom & Ho, *supra* note 6, at 41.

113 Recent hacking and data breaches in government agencies have sparked greater interest in technology supply-chain risks, reinforcing the critical importance of documenting provenance at each step in the development of an algorithmic tool, even when the work is done by "trusted" stakeholders.

114 *E.g.*, Engstrom & Ho, *supra* note 25, at 6-7; Doshi-Velez et al., *supra* note 7.

115 Shapiro et al., *supra* note 106, at 493-94.

116 *Id.*

117 Brian D. Feinstein & Daniel J. Hemel, *Outside Advisers Inside Agencies*, 108 GEO. L. J. 1139 (2020); McGarity & Wagner, *supra* note 77, at 1800-01.

118 Similar concerns might arise in the selection of internal scientists as peer reviewers, a topic that lies outside the scope of this study.

119 *But see* ACUS AI Study, *supra* note 29, at 18 (53% of the AI applications in the study were built in-house by government).

120 Government contractors might best be treated as private parties with respect to the creation of AI tools, but we simply raise this issue and do not address it here.

121 *See generally* ACUS AI Study, *supra* note 29.

122 *E.g.*, The A.I Ctr. of Excellence, *How We Measure AI Readiness*, IT Modernization Ctr.'s of Excellence (Oct. 28, 2020), https://coe.gsa.gov/2020/10/28/ai-update-2.html [https://perma.cc/5PKF-YCHB].

About the Authors

Wendy Wagner is the Richard Dale Endowed Chair at the University of Texas School of Law. Wagner earned her law degree from Yale Law School and a masters in environmental studies also from Yale, clerked for the Honorable Albert Engle of the U.S. Sixth Circuit Court of Appeals, and served as an honors attorney in the environmental enforcement section of the Department of Justice before joining Case Law School (her first academic appointment) in 1992. Wagner's research focuses on issues related to the design of bureaucratic processes, environmental regulation, and law and science. She has authored three books and dozens of articles. Outside of her academic duties, Wagner has participated in various organizations, including the American Bar Association, National Academies of Science committees, the Bipartisan Policy Center, the Center for Progressive Reform, and served as a consultant to the Administrative Conference of the U.S.

Martin Murillo is a researcher focused on the intersection of political science and engineering. His work includes early conceptualization of the role of data and algorithms in government transparency, the role of data breaches on government transparency and accountability, and the addressing of jurisdictional issues of Internet of Things architectures and technologies. His peer-reviewed work has been published in venues such as the International Political Science Review, the IEEE Internet of Things Journal, the American Control Conference, and the American Institute of Aeronautics and Astronautics. He has participated in policy-related venues sponsored by the W3C, IEEE, UNFCCC, and IETF with the aim to influence the development of more inclusive standards and technologies. He holds a B.S. degree from Brigham Young University and a Ph.D. degree from Idaho State University; he has also done post-doctoral work at the Norwegian Institute of Science and Technology. In addition to his current research on algorithms and governance, he manages an IEEE-sponsored project that utilizes grassroots initiatives for data connectivity of remote underserved politically autonomous communities.

© 2021, Wendy Wagner and Martin Murillo.

Acknowledgments

The authors thank Katy Glenn Bass, Adam Glenn, Jameel Jaffer, Amy Kapczynzski, and the Knight Institute staff for the opportunity to work on this project and for excellent suggestions along the way.

About the Knight First Amendment Institute

The Knight First Amendment Institute at Columbia University defends the freedoms of speech and the press in the digital age through strategic litigation, research, and public education. It promotes a system of free expression that is open and inclusive, that broadens and elevates public discourse, and that fosters creativity, accountability, and effective self-government.

knightcolumbia.org

Design: Point Five
Illustration: ©Erik Carter



KNIGHT
FIRST AMENDMENT
INSTITUTE at
COLUMBIA UNIVERSITY