Legitimacy Practices Under Algorithmic Governance

by

Amina A. Abdu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Information)
in the University of Michigan
2026

Doctoral Committee:

Professor Abigail Jacobs, Chair
Professor Elizabeth Popp Berman
Professor Ben Green
Professor Simone Zhang

Amina A. Abdu
aabdu@umich.edu
ORCID iD: 0009-0008-9505-439X

© Amina A. Abdu 2026

ACKNOWLEDGEMENTS

I have many people to thank for helping to make this work possible. First and foremost, thank you to my advisor Abigail Jacobs, who has been an invaluable collaborator on the work presented in this dissertation, and an even more invaluable mentor throughout my PhD. Thank you for teaching me about the importance research taste and giving me both the space and the support to develop it. I deeply appreciate the guidance of my other dissertation committee members: Elizabeth Popp Berman, Ben Green, and Simone Zhang. I have been so lucky to learn from our conversations and hope I will get to keep doing it for a long time yet. Thank you also to the many people who have provided helpful suggestions on various parts of this work. In particular, I would like to acknowledge participants at the Keywords of the Datafied State workshop, 2024 Privacy Law Scholars Conference, the 2024 Law and Society Annual Meeting, the 2024 Privacy and Public Policy Conference, and the 2024–2025 Rackham Interdisciplinary Workshop in STS. I am also grateful for funding from the U-M Initiative for Democracy & Civic Empowerment. I am fortunate to have benefited from both intellectual community and friendships at UMSI, the U-M STS community, FAccT, and PLSC. All of these spaces and the people in them have given me a place to call my academic home at times when I felt unmoored by my work. I would also like to thank the members of my pre-lim committee– Irene Pasquetto and Joy Rohde– who helped shape many of the early ideas in this dissertation and helped me understand who I was (and could be) as a researcher. I am lucky to have gotten incredible feedback and support over the years from my current and former labmates: Meera Desai, Hana Galijasevic, Jared Katzman, and Lu Xian. A huge thank you to my collaborators: to Lauren Chambers and Deirdre Mulligan, who first got me excited about technology in federal agencies, and to Jeremy Seeman, who has been a wonderful collaborator on the work presented in Chapter 4. It is impossible to express how grateful I am for my friends and family near and far, whose steady presence has been a welcome reminder of life outside of this dissertation. Christian, I could not have done this without your support– which is miraculously as formidable from 800 miles away as it is when we get to be in the same room. How lucky am I to have my biggest champion and my favorite thinker in the same person? Finally, none of this would have been possible without my parents, Iman & Ashraf, whose love, support, and lifelong belief in me have sustained me through so many chapters and have made this the most fulfilling one yet.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS ii

LIST OF FIGURES v

LIST OF ACRONYMS vi

ABSTRACT vii

CHAPTER

1 Introduction **1**

2 Background **6**

2.1 A Note on Legitimacy 6
2.2 Legitimacy in the U.S. Administrative State 7
2.3 From the Administrative State to the Algorithmic State 9
2.3.1 Algorithms and arbitrariness 9
2.3.2 Algorithms and accountability 10
2.4 The Legacy of Quantification 12

3 Legitimacy Practices, Algorithms, and the New Bureaucratic Quantification **14**

3.1 Introduction 14
3.1.1 Roadmap 15
3.2 Legitimacy Practices 16
3.3 Administrative Legitimacy and Quantification 17
3.3.1 Quantification and Efficiency 17
3.3.2 Quantification and Non-Arbitrariness 18
3.3.3 Quantification and Accountability 19
3.4 Methods 19
3.4.1 Study Setting 19
3.4.2 Data 20
3.4.3 Empirical Strategy 21
3.5 Findings 23
3.5.1 Findings: Efficiency in the Algorithmic Administrative State 23
3.5.2 Findings: Non-Arbitrariness in the Algorithmic Administrative State 26
3.5.3 Findings: Accountability in the Algorithmic Administrative State . 29

3.6 Discussion . . . 31
3.6.1 Limitations . . . 34
3.7 Conclusion . . . 35
4 Ethical Closure Problems in Formalizing Technologies . . . 36
4.1 Introduction: "Perfect Technologies of Justice" . . . 36
4.2 Formalizing Technologies . . . 38
4.2.1 Technologies of Quantification . . . 38
4.2.2 Formalizing Technologies . . . 39
4.3 From Scientific to Ethical Closure . . . 41
4.3.1 Scientific Closure . . . 41
4.3.2 Ethical Closure . . . 42
4.4 Closure in Practice . . . 43
4.4.1 Case 1: Formal Privacy-Enhancing Technologies . . . 43
4.4.2 Case 2: Game Theoretic Decision-Making and Mechanism Design . . . 45
4.5 FTs as Rhetorical Technologies . . . 48
4.5.1 Rhetorical Technologies . . . 48
4.5.2 Beyond Law, Beyond Quantification . . . 51
4.6 Conclusion . . . 54
5 Error, Fraud, and Fodder for the Accountability Machine: The Negotiation of Administrative Error Metrics in SNAP . . . 56
5.1 Introduction . . . 56
5.2 Background . . . 58
5.3 Accountability Metrics in Practice: Supplemental Nutrition Assistance Program (SNAP) Quality Control . . . 62
5.4 From Metrics to Accountability: a Typology of Levers . . . 64
5.4.1 Operationalization . . . 65
5.4.2 Incentives . . . 67
5.4.3 Auditing System . . . 71
5.4.4 Interactions Between Levers . . . 72
5.5 Discussion . . . 73
5.6 Conclusion . . . 75
6 Conclusion . . . 77
6.1 Reflections and Future Directions . . . 78
BIBLIOGRAPHY . . . 80

LIST OF FIGURES

FIGURE

3.1 Illustrative examples of GAO reports in our sample. Political anxieties can direct attention to agency activities to target different concerns (e.g., fraud, unemployment, national security) and perform different values (e.g., efficiency, flexibility, consistency). 20

3.2 GAO reports 1968-2024. Top: Histogram of the number of GAO reports including the terms "algorithm" and/or "artificial intelligence" per year. (Note: our data collection stopped March 18, 2024, but this table includes all reports through the end of 2024.) Bottom: Histogram of all GAO reports during the same period for comparison. 22

3.3 **Themes of efficiency** as a value emphasize the promises of algorithmic adoption across a range of expressions of efficiency, but also appear as a potential challenge of algorithmic adoption and in interventions for establishing legitimacy through other means. 24

3.4 **Themes of non-arbitrariness** are called upon to address the promises, challenges, and legitimacy practices around the adoption of algorithmic technologies. 27

3.5 **Themes of accountability** in the GAO reports emphasize the challenges of algorithmic adoption and how legitimacy practices could achieve accountability via documentation. 30

LIST OF ACRONYMS

AI artificial intelligence

APA Administrative Procedure Act

CAPER Case and Procedural Error Rate

DOGE Department of Government Efficiency

DP differential privacy

FNS United States Department of Agriculture Food and Nutrition Service

FP formal privacy

FT formalizing technology

FTC Federal Trade Commission

GAO Government Accountability Office

ICC Interstate Commerce Commission

NSF National Science Foundation

PER Payment Error Rate

PET privacy-enhancing technology

PLB privacy loss budget

QC Quality Control

SNAP Supplemental Nutrition Assistance Program

STS science and technology studies

USDA United States Department of Agriculture

ABSTRACT

This dissertation examines how algorithmic technologies are transforming the everyday institutional practice and performance of legitimacy in U.S. public administration. Bridging theories of legitimacy from organizational sociology, science and technology studies (STS), and administrative law, I introduce *legitimacy practices* as a crucial site where the scope of legitimate policy interventions is negotiated, focusing in particular on algorithmic governance – both in the sense of policy implemented by algorithm and regulation of algorithms. Prevailing approaches treat legitimacy as a static construct, which is either undermined or furthered when agency officials turn to algorithms. However, this dissertation proposes that the algorithmic turn within government agencies reflects a change in how administrative legitimacy is understood and performed, with value-laden consequences. I examine this algorithmic turn across three empirical contexts. First, I chart the Government Accountability Office's role in managing administrative legitimacy through its everyday oversight of agency activities. I discuss how algorithms have reconfigured these practices, elevating efficiency and accuracy, while sidelining accountability. Second, I explore the rhetorical strength of technologies that attempt to formalize ethical problems – including formal privacy and game theory – to understand how different forms of algorithmic public administration make and are made legitimate by scoping the space of policy contestation. Finally, I examine how data infrastructures institutionalize agency accountability through an examination of Supplemental Nutrition Assistance Program (SNAP) error rates, which govern relationships between federal and local governments, administrative agencies and the executive branch, and the state and the public. Through these distinct contexts, I show how the push toward algorithms reflects a turn toward efficiency and accuracy as the core legitimating principles of the administrative state — and a concurrent side-lining of reason-giving and democratic accountability. Together, these three studies help situate the algorithmic turn within the longer history of quantification in public administration, revealing changes to who (or what) is given the authority to implement policy and on what grounds.

CHAPTER 1

Introduction

Algorithmic technologies, including artificial intelligence (AI), increasingly inform high-stakes policy decisions within government, automating administrative tasks from privacy protection to fraud detection. Despite its rapid adoption, we don't yet understand how algorithmic technologies relocate political power and threaten the values underpinning democratic processes. This dissertation intervenes in this space by surfacing the social implications of technical implementation decisions and the processes by which these decisions are made.

In this dissertation, I examine how algorithmic technologies and technical experts shape policy through changing ideas of legitimacy. I address these questions from multiple perspectives, examining how governments, academic communities, and civil society engage with algorithmic technologies during policy-making processes. I argue that, in quantitatively operationalizing how policy is enacted, algorithms play a key role of public administration, taking on a role traditionally filled by agency officials. The authority of these unelected civil servants to make policy decisions has long been contested and the turn to algorithms has reanimated many familiar anxieties about the legitimacy of the administrative state. Questions about legitimacy govern who has the ability to make policy and on what grounds. In the context of public administration, this affects how public resources are allocated and the public's ability to have a say in how they are governed across areas like healthcare, nutrition, military intervention, and the environment. The introduction of algorithms into these contexts threatens to alter the relationship between the public and the state, transforming access to essential social services, circumventing established pathways for democratic deliberation, and undermining the role of the administrative state altogether. As state and federal agencies face sustained political attacks and dwindling resources, the adoption of algorithms does not just replace the work done by agency officials, but alters what work agencies have the authority to do and why.

In this dissertation, I show that the adoption of algorithms in agencies is both justified by upstream legitimacy concerns and creates downstream consequences for legitimacy. Legal and philosophical approaches to legitimacy typically focus on whether authority is justified

based on existing legal or moral frameworks. Meanwhile, sociological approaches often focus on whether the subjects of this authority perceive it as justified. My work complements these existing approaches by instead examining how legitimacy is performed by those in power for other authorities. This conception of legitimacy generates insights into how power and authority are negotiated over time and, in the case of this dissertation, when new technologies are introduced into the political process. This analysis draws on and brings together theories of legitimacy from administrative law, organizational theory, and science and technology studies (STS). Focusing on the U.S. government's use of algorithms in state and federal agencies, I examine how different actors' everyday algorithmic practices reconfigure what constitutes a legitimate policy intervention.

This dissertation brings together three studies to examine how the use of algorithmic technologies reconfigure the practices that agencies use to demonstrate their legitimacy. The term "algorithm" is used to describe a wide range of technologies—from automated rule-based systems to general-purpose large language models— often emphasizing varied technical features including automation, "black-box" inscrutability, stochasticity, and data-driven-ness. Rather than prescribing a particular technical definition of algorithms, this dissertation focuses on what actors themselves refer to as algorithms. This approach treats the term "algorithm" not merely as a technical category, but also as a legitimacy practice: a term that actors use to describe and justify particular technologies and interventions. Attending to how actors deploy the term makes it possible to examine how the classification of something as an algorithm itself participates in the production of legitimacy.

In contexts spanning internal government management, individual data privacy, warfare, and access to food, these studies demonstrate how a broad range of algorithmic technologies alter what comes to be seen as a legitimate or illegitimate exercise of government authority. In particular, this dissertation addresses the following questions: When and why do agencies turn to algorithmic technologies in pursuit of legitimacy? How do algorithmic technologies relocate power in policy-making processes? And, in parallel, how do algorithmic technologies reshape the types of knowledge, expertise, and values that are seen as legitimate foundations for policy interventions? I answer these questions across three sites of study– the Government Accountability Office (GAO)'s oversight of agency algorithm use, technologies that attempt to formalize the relationship between policy inputs and social goals, and the data infrastructures underlying fraud detection in the state-level administration of the SNAP program. These sites comprise multiple actors– government officials, academics, technologists, and citizens– that enact distinct legitimacy practices through their use of and rhetoric about algorithmic technologies. Leveraging the rhetoric and policy outputs of each of these groups, I reveal how these actors negotiate and redefine legitimacy, expertise, participation,

and, ultimately, the political process.

The rest of the dissertation proceeds as follows. In Chapter 2, I provide background on the U.S. administrative state's efforts to establish its own legitimacy, and review the literature on algorithms and legitimacy within this context. In Chapters 3–5, I include three studies that address legitimacy practices surrounding algorithmic technologies before concluding in Chapter 6. Each of the following subsections previews one of these three studies.

Chapter 3: Legitimacy Practices

Chapter 3 asks how algorithmic technologies are transforming the everyday institutional practice and performance of legitimacy, focusing on U.S. public administration. This is both a high-stakes domain where algorithms are deployed, and a site where the legitimacy of algorithmic technologies is particularly salient because of long-standing anxieties about the legitimacy of agency authority within the administrative state. These anxieties have meaningfully shaped agency practices, which have sought to produce legitimacy by minimizing arbitrariness, expanding accountability, and enhancing efficiency. In turn, agency practices have shaped the meaning of legitimacy, uplifting certain values and bureaucratic procedures as necessary means to demonstrate legitimacy. Quantification and the turn to state statistics have been particularly important in this history of establishing agency legitimacy. While the turn to algorithms in government similarly relies on quantification, its relationship to legitimacy is less clear. This chapter examines the turn to algorithms as an emerging form of quantification, with particular implications for administrative legitimacy.

I argue that, in quantitatively operationalizing how policy is enacted, government algorithms play a key role of public administration (Abdu and Jacobs, 2024), taking on a role traditionally filled by agency officials. The turn to algorithms has reanimated many familiar anxieties about the legitimacy of the administrative state. This chapter introduces everyday legitimacy practices as a crucial site where the meaning of administrative legitimacy is negotiated, focusing in particular on practices around the use of algorithms within federal agencies. The U.S. Government Accountability Office (GAO) acts squarely in the context of everyday state legitimacy practices. GAO is an independent agency in the legislative branch that provides oversight over the federal government, including administrative agencies in the executive branch, on behalf of Congress. GAO manages the legitimacy of administrative agencies through its everyday oversight, emphasizing nonpartisan information gathering and accountability. Drawing on GAO reports, I empirically examine how algorithms can reshape the values, goals, and logics of legitimacy, situating the algorithmic turn within the long history of quantification. This chapter shows that the turn to algorithms extends the logic

of quantification– toward efficiency and accuracy as legitimating principles– but has not been granted the same legitimacy as prior forms of quantification. Instead, algorithms have until recently been deployed primarily in internal rather than public-facing settings. This precarious legitimacy combined with the continued emphasis on efficiency at all costs places agencies in a no-win situation in efforts to justify their authority.

Chapter 4: The Rhetoric of Technology

In Chapter 4, I examine the rhetorical power of technologies that mathematically formalize policy outcomes. Using the framework of scientific closure, I outline how formalizing technologies produce a particularly scientific notion of legitimacy, narrowing the space of ethical decisions in bureaucratic contexts. I describe how this process of closure occurs and its impact on shaping regulatory decision-making and practice. This chapter draws on two case studies: the U.S. Census Bureau's adoption of differential privacy, and military applications of game theory. I use these two case studies from different times and across different domains to show how the adoption of formalizing technologies reconfigured the space of legitimate interventions and foreclosed certain forms of accountability and participation. Through this work I argue for a rhetorical approach to technology which examines how particular kinds of technologies frame policy problems. Distinguishing formal quantification from other forms of quantification reveals how specific technological affordances reshape who is seen as a legitimate participant in the policy-making process, and where they can intervene. In the case of formalizing technologies, their emphasis on ex-ante guarantees rather than empirical data forecloses contestation and participation based on evidence, instead limiting participation to upstream discussions about assumptions that have often occurred in less visible, less democratic spaces.

Chapter 5: Shifting and Conflicting Accountabilities

In Chapter 5, I examine how accountability is enacted through data infrastructures. The first two chapters of my dissertation describe how many forms of accountability have been sidelined in the performance of administrative legitimacy. This chapter seeks to understand and chart the evolution of administrative accountability and its relationship to legitimacy more closely. In this study, I examine how quantification shapes accountability, highlighting the different components of quantification that meaningfully alter administrative accountability. As the prior chapters highlight, accuracy has become a central tenet of legitimacy under algorithmic governance. Error metrics, therefore, promise to identify and eliminate sources of

illegitimacy. Error rates are used to justify the adoption of algorithms to root out fraud and waste, and provide the inputs upon which these algorithmic systems are built. These metrics act as a crucial infrastructure underpinning administrative accountability under algorithmic governance. This chapter endeavors to understand how changes to these metrics reshape accountability within the administrative state, with important consequences for access to public services and individual privacy.

I focus in particular on the SNAP Quality Control (QC) system, which involves multiple metrics constructed based on state records of SNAP payments to verify the "integrity" of social benefits administration. I examine the wide range of accountability relationships mediated by QC metrics: between state and federal agencies, between administrative agencies and the executive branch, and between the government and civil society. I discuss how technical changes to the QC system affect how these accountability relationships are constructed. Analyzing the QC system through an infrastructural lens, this chapter demonstrate how accountability practices can become embedded within government even when legitimacy practices have otherwise been subverted to minimize the role of accountability. By examining variations over time, across states, and across various kinds of accountability relationships, I argue that this site illuminates where and for whom accountability still acts as a meaningful part of legitimacy practices. Focusing on federal and state-level policy changes around QC error rates, I identify three components of metrics that meaningfully shape accountability: how metrics are *operationalized*, the *incentives* tied to metrics, and the broader *auditing system* that determines how these error rates are constructed on the ground. A closer look at these components reveal that accountability has shifted to de-emphasize the public's role in administrative accountability but that this need not be the case.

CHAPTER 2

Background:

Legitimacy in the Algorithmic Administrative State[1]

This dissertation contributes to the growing literature on the legitimacy of algorithmic governance (in the sense of government-by-algorithm), particularly in the context of public administration. As administrative tasks are turned over to algorithmic systems, many longstanding concerns about the legitimacy of the administrative state have been cast into a new light. In this chapter, I outline these longstanding concerns, and review some common arguments about how agency adoption of contemporary algorithmic systems to administer public policy challenge the historical means by which the administrative state has justified its legitimacy– minimizing arbitrariness and instituting accountability. Finally, I lay out the need to re-examine administrative legitimacy practices to understand how agencies are reinterpreting legitimacy in light of increasing algorithmic adoption. While these new interpretations of legitimacy are not necessarily caused by the turn to algorithms, the deployment and justification of algorithmic technology can facilitate changing norms around who has the power to make policy decisions, and on what grounds.

2.1 A Note on Legitimacy

I define legitimacy as the moral justification for exercising authority. I take a discursive approach to understanding legitimacy, examining how those in power justify their own exercise of authority or how oversight institutions justify the authority of others. Specifically I focus on *administrative legitimacy*, situating agency practices and discourses around algorithms within the context of the U.S. administrative state's long-standing efforts to justify

[1]Parts of this chapter are modified from the following essay previously published as part of Data & Society's *Keywords of the Datafied State*: **Amina Abdu** & Abigail Jacobs. "Public Administration." In Burrell, Jenna, Ranjit Singh, and Patrick Davison (eds.), *Keywords of the Datafied State*. Data & Society. 2024.

its authority. Prevailing approaches treat legitimacy as a static construct, which is either undermined or furthered when agency officials turn to algorithms. Legal and philosophical approaches to legitimacy typically focus on whether authority is justified based on existing legal or moral frameworks (Danaher, 2016; Calo and Citron, 2020; Lazar, 2024; Stone and Mittelstadt, 2024). Meanwhile, sociological approaches often focus on whether the subjects of this authority perceive it as justified (Tyler, 1990; Waldman and Martin, 2022; Martin and Waldman, 2023). Generalizing across a wide range of sociological approaches to legitimacy, Shoon (2022) proposes a dyadic conception of legitimacy which requires an *object* of legitimacy, an *audience* which confers legitimacy, and a relationship between the two that is used to assess legitimacy through *expectations*, *assent*, and *conformity to expectations*. In this dissertation, I take a particular approach to the question of the audience: examining how legitimacy is performed by those in power for other authorities and expert communities. This conception of legitimacy generates insights into how power, expertise, and authority are negotiated over time. As we will see, often these negotiations bypass the subjects of authority who ultimately bear the consequences for what is seen as legitimate. By focusing on *legitimacy practices*, this dissertation shows how the expectations underpinning legitimacy are reconfigured when new technologies are introduced into the political process, in turn changing what it means to exercise legitimate power.

2.2 Legitimacy in the U.S. Administrative State

The Federal government in the United States comprises elected representatives, like the president and members of Congress, and the federal agencies that do the work of implementing public policy, also known as public administration. Unlike members of Congress, agency officials are appointed rather than elected to work in agencies such as the Food and Drug Administration, the Environmental Protection Agency, or the Federal Trade Commission. In the early 20th century, public administration in the U.S. was envisioned as the apolitical implementation of the state's will, as expressed by the legislative branch. However, the work of agency officials has since come to be understood as political work in its own right. Thus, at the core of public administration's history in the United States, is an ongoing contest over the legitimacy of the administrative state: while the authority of legislators to make political decisions comes from democratic elections, agency officials must look elsewhere. Yet to make and administer policy, government agencies must depend on the trust they foster. The efficiency and efficacy of agencies thus depends on their perceived legitimacy. The U.S. government's strategies to legitimate agency authority has repeatedly shaped the meaning of public administration and, conversely, changes in public administration have

raised new questions of legitimacy. Recently, public administration has once again taken a new shape as agency officials turn to algorithmic systems to automate parts of the policy implementation process previously performed by agency experts. Agencies tout the promise of more objective and efficient government administration — while algorithmic systems also present risks of bias, errors, and opacity — all with uncertain consequences for legitimacy. We examine the turn to algorithms in public administration within the longer history of U.S. agencies' efforts to pursue legitimacy. Two strategies in particular have been central to the American administrative state's pursuit of legitimacy: minimizing arbitrariness and expanding accountability (Bressman, 2003). This historical struggle is newly salient in the algorithmic state, in which the locus of administration is moving from government officials to data-driven technologies raising familiar issues of legitimacy, efficiency, arbitrariness, and accountability. While the turn to algorithms can be seen as one in a long history of efforts to strengthen the legitimacy of the administrative state, it also runs the risk of undermining this hard-fought legitimacy.

In 1887, the Interstate Commerce Commission (ICC) was established to regulate abuse of monopoly power in the railroad industry. The ICC was the first independent federal agency, and in many ways, the beginning of the administrative state, as the U.S. government subsequently established other permanent agencies to take on the work of public administration. As the administrative state grew in the early 20th century, legal scholars began to debate the legitimacy of these agencies under the U.S. Constitution. As Lisa Bressman (2003) describes, early paradigms of public administration were primarily concerned with the threat of arbitrary decisions — decisions based on agency officials' individual whims and biases rather than reasoning. The administrative state initially pursued legitimacy through efforts to minimize arbitrariness in the implementation and execution of the "objective" policies dictated by Congress. The federal government sought to mitigate the threat of arbitrariness first through detailed legislative directives and then, when this proved inefficient during the New Deal expansion of the federal government, through appeals to scientific objectivity and the hiring of scientific experts. Finally, in 1946, Congress enacted the Administrative Procedure Act (APA) to standardize agency procedures. These strategies sought to constrain agency discretion in order to produce consistent, predictable, and fair decisions. By the 1970s, however, idealism about objectivity waned and the administrative state shifted its focus from minimizing arbitrariness and discretion to expanding agency accountability.

Accountability requires not just transparency but a richer relationship that Mark Bovens (2007) defines as "a relationship between an actor and a forum, in which the actor has the obligation to explain and justify his or her conduct, the forum can pose questions and pass judgment, and the actor might face consequences." The administrative state sought

accountability through multiple channels. In the 1970s and 80s, lower courts and, eventually, the U.S. Supreme Court began to require that federal agencies provide justifications for their decisions. This not only enabled greater public accountability through transparency requirements, but also encouraged consistent decision-making by demanding reasoning around deviations from prior judgments. Agencies pursued legitimacy through expanded public participation, like "notice and comment" rulemaking. Defined in the APA notice-and-comment requires that agencies notify the public of proposed rules and allow the public to submit comments before issuing a final rule. In theory, notice-and-comment relied on the same method as elected officials: ideals of democratic representation and inclusion. However, public participation proved challenging. Reading and responding to comments imposed an efficiency burden on agencies and many questioned whether certain groups were being systematically over- or under-represented in participatory political processes. Gradually, beginning with the Reagan Administration, the administrative state has pursued accountability by bringing the administrative apparatus under the control of the president, who is subject to the consequences of public opinion. These three mechanisms — transparency requirements, public participation, and the possibility of consequences — reflect key dimensions of accountability, each of which have been used to support the legitimacy of the American administrative state.

2.3 From the Administrative State to the Algorithmic State

Under algorithmic governance, the administration of policy is increasingly tasked to automated decision-making systems. Such algorithmic technologies are frequently introduced and justified as a way to ensure more efficient, more accurate, and less arbitrary decision-making (Engstrom et al., 2020). However, critical scholars note that these technologies pose a significant threat to accountability, transparency, contestability, and public participation. Now with algorithmic systems as the focus, the fight for the legitimacy of the administrative state is being rehashed, with algorithmic systems being used to pursue familiar goals of efficiency, arbitrariness, and accountability but also raising novel challenges, such as the opacity of algorithmic systems, the reduced role of deliberation and reasoning, and the displacement of agency expertise.

2.3.1 Algorithms and arbitrariness

Discussions around government algorithms have centered on many of the same ideals as those pursued by the administrative state: legitimacy, once again contested on the grounds

of arbitrariness, accountability, and efficiency. Proponents of administrative algorithms highlight their potential to advance legitimacy, often pointing to increased efficiency and reduced arbitrariness. Critics, however, contend that algorithmic systems in fact pose unique challenges to non-arbitrariness, while simultaneously undermining accountability. Echoing early administrative goals with appeals to efficiency, Madalina Busuioc (2021) describes how "the promise of efficient, low-cost, or 'neutral' solutions harnessing the potential of big data has led public bodies to adopt algorithmic systems in the provision of public services." In addition to promises of improved efficiency, proposed benefits of algorithmic technology in government include reduced bias, increased accuracy, and increased consistency. In theory, algorithmic technology has the potential to reduce inconsistencies and the place of individual discretion or whims of agency officials through the implementation of rules. In other words, government algorithms promise to mitigate the arbitrariness of decisions made within the federal agencies that use them. Indeed, central to the promise of algorithmic decision-making is the veneer of "objectivity" and political neutrality pursued throughout the administrative state's long-standing efforts to establish its own legitimacy. In particular, the turn toward automated systems reflects a push toward what Lorraine Daston and Peter Galison (2007) call mechanical objectivity, a vision of objectivity that suppresses human judgment in favor of rules. Prizing mechanical objectivity reflects a vision of administration where arbitrariness is seen as driven by human problems of bias, error, and inconsistency — rather than a lack of reasoning or moral justification.

However, critics of government algorithms and algorithms used in other high-stakes settings have highlighted that algorithmic decision-making systems both reproduce and exacerbate the human bias it purports to solve as well as simply make mistakes. Randomness is necessarily built into machine learning models, which presents a legal obstacle to non-arbitrariness (Cooper et al., 2022a). Moreover, objectivity is neither achieved nor achievable: systems used to administer policy necessarily quantify social goals and encode values. Algorithmic systems, like human administrators before them, can never implement policy without themselves making policy decisions (Mulligan and Bamberger, 2018). Nor can they fully eliminate values or arbitrariness from these policy decisions. Thus, like previous iterations of the administrative state, their legitimacy must rest on their accountability to the public.

2.3.2 Algorithms and accountability

Algorithmic technologies can reshape accountability in important ways. The quantification implicit in algorithms promise to promote accountability through transparency and public

participation alike. Quantification in government, such as cost-benefit analysis, promises to make political decisions more visible through the communication of verifiable numbers. Meanwhile, quantification as an instrument of dispute enables stakeholders and publics to participate in democratic debate and challenge powerful institutions in a common and authoritative language. Perhaps as a result of these features, accountability, inside and outside of government, has increasingly been conceptualized and implemented in quantitative terms, often perceived as transparent, objective, and universal.

Discussions around accountable technologies in administration have focused on transparency as a means to accountability. Yet accountability is not solved by transparency alone nor explainability. Contestability mechanisms that allow the public to meaningfully shape and challenge systems must be intentionally designed. Previously domain experts, policymakers, and publics were involved in developing and executing policies, where technologies are increasingly taking this role. However, that expertise is now overlooked in these discussions. The development of technologies does not eliminate judgment, but instead threatens to displace it to less visible and less accountable locations, where software is developed away from experts and its developers are, unlike traditional government officials, unaccountable to the public.

Despite the promise of quantification, the lack of accountability afforded by algorithmic decision-making systems is often cited as a primary challenge for their adoption in government. The algorithmic accountability literature has focused in particular on the lack of transparency (Citron, 2007; Kaminski, 2020) and explainability (Ajunwa, 2023; Busuioc, 2021) of automated decision-making systems. The opacity and inscrutability of complex or otherwise hidden algorithmic systems challenges accountability by minimizing the possibility of informed participation (Vaccaro et al., 2019). A. Feder Cooper et al. (2022b) identify a number of obstacles to algorithmic accountability surrounding the attribution of moral responsibility and subsequent susceptibility to consequences. The number and variety of people involved in the machine learning pipeline creates a problem of "many hands," obscuring both who is responsible for any given decision within a complex system as well as to whom accountability is owed. Even when the division of labor within a system is clear, "bugs," or issues that are seen as inherent in the coding process, make it difficult to attribute a cause to a given problem while computers provide an un-sanctionable scapegoat for human and institutional decisions. The government procurement process by which agencies acquire algorithmic systems poses a particular challenge to government accountability because agency experts abdicate responsibility for policy decisions to third-party developers who are not subject to the same public participation, deliberation, or oversight (Jacobs and Mulligan, 2022). Compounding these issues, private technology companies, many of whom

are involved in the administrative procurement process, have successfully managed to minimize legal liability further minimizing the possibility of sanctions, one of the key elements of accountability (Cooper et al., 2022b).

By challenging justification, participation, and consequences, algorithms challenge many of the logics central to administrative accountability. As traditional notions of accountability, and even non-arbitrariness, are under strained, this dissertation seeks to understand how public administrators in the U.S. are reconfiguring their legitimacy practices to justify government-by-algorithm.

2.4 The Legacy of Quantification

In this dissertation, I situate the turn to government algorithms within the proliferation of quantification in agency practices. From state censuses to cost-benefit analysis, the use of numbers in administrative contexts has served as an important tool for establishing and maintaining the legitimacy of agency authority (Porter, 1992). Scholars of quantification have examined the close link between administration and numbers, focusing on the capacity to assess, predict, and persuade (Mennicken and Espeland, 2019). Quantification practices align with the administrative state's pursuit of legitimacy through promises of non-arbitrariness and accountability. Quantification reflects an attempt to minimize arbitrariness through recourse to objectivity, as well as through appeals to consistency and fairness. Simultaneously, quantification endeavors to expand accountability by producing trust, by allowing actors across contexts to raise and settle disputes, and by enabling visibility (Espeland and Vannebo, 2007; Porter, 1995; Deringer, 2018; Igo, 2007).

The accountability, and non-arbitrariness promises of quantification have given numbers an important place among the procedures used to demonstrate agency legitimacy. Algorithmic technologies, which build on quantitative tools to automate analysis and classification decisions, have increasingly been proposed and implemented in agency workflows (Engstrom et al., 2020). Yet, algorithms do not necessarily fall neatly among other technologies of quantification: in their review of work on the sociology of quantification, Berman and Hirschman (2018) ask: "Are algorithms quantification?" (p. 258). This dissertation traces continuities between quantification and algorithms, while highlighting the usefulness of distinguishing between different forms of quantification by attending to specific technical features of algorithms, and specific rhetoric around algorithms.

The literature on algorithmic legitimacy has focused on how algorithms support—or more often, and contrary to other forms of quantification, undermine—administrative legitimacy. However, notions of legitimacy are not stable, but change with practice. As the quantification

literature makes clear, consequential political choices about how numbers are constructed and deployed often fade to the background once these numbers become naturalized (Barman, 2016; Merry, 2016). Focusing on algorithms, which are in an ongoing process of naturalization, we examine the process by which these choices are naturalized through negotiations around technology, expertise, political values, and legitimacy. Quantification is not only used to pursue legitimacy: it also changes how legitimacy is enacted. The proliferation of quantification in government has redefined legitimacy practices within policy spaces, facilitating demonstrations of efficiency and elevating efficiency as a key legitimating principle of good government (Demortain, 2019): "Quantification makes individual and organizational performance visible, trackable, and comparable, thereby allowing for organizing in accordance with principles of efficiency"(Mennicken and Espeland, 2019, p. 233). In this dissertation, I show that efficiency is often used to justify the adoption of algorithms, even when it undermines traditional legitimacy values like accountability and non-arbitrariness, extending a core legacy of quantification and highlighting the need to examine how algorithms support changing notions of legitimacy. In addition to treating algorithms broadly (Ch. 3), I distinguish between different forms of quantification that underpin algorithms including formal modeling (Ch. 4) and metrics (Ch. 5). These categories are not intended to provide an exhaustive taxonomy of quantification, but to highlight distinctions between different technologies that fall under the umbrella of quantification. While algorithms share many common features with other forms of quantification, attention to particular technological affordances can illuminate where new technologies are deployed toward changing visions of legitimacy.

CHAPTER 3

Legitimacy Practices, Algorithms, and the New Bureaucratic Quantification[1]

3.1 Introduction

In the United States, experts within administrative agencies play a critical role in overseeing complex social problems and responding quickly to emerging issues. Yet the authority of agency policymakers—who are appointed rather than elected—and the legitimacy of the administrative state more fundamentally have been the subject of debate for as long as these agencies have existed (Bressman, 2003; Beermann, 2024; Fisch, 2024). Early models of public administration located the legitimacy—and value-neutrality (Seidenfeld, 1994)—of agency decision-making in apolitical scientific expertise and fidelity to congressional statutes (Seifter, 2014). However, the interpretation and implementation of policy requires managing ambiguities inherent to both law and science (Kumar, 2025; Larkin, 2020). Whether agency officials have the authority to resolve such ambiguities, and how those ambiguities are resolved, has thus long been a locus of political tension.

The ambiguity of public administration invites quantification. Numbers promise impartial and rational adjudication of difficult policy questions. Historically, agencies have emphasized quantification to assert legitimacy in contexts where trust and consensus are difficult to come by (Porter, 1995). New forms of quantification, however, challenge this logic. Specifically, the adoption of algorithms within agency operations is reanimating many long-held tensions in the legitimacy of the administrative state (Abdu and Jacobs, 2024). Shifting authority from bureaucrats to algorithms (and the technologists who design them) raises questions as to whether and how this authority is legitimized (Hong, 2023). Despite substantively overlapping concerns and mechanisms, the literature on quantification in public administration

[1]This chapter is drawn from the following paper: **Amina A. Abdu** and Abigail Z. Jacobs. (2026). Legitimacy practices, algorithms, and the new bureaucratic quantification. *Big Data & Society, 13*(2)."

has been largely disconnected from the literature on algorithmic legitimacy and public administration. In this article, we propose that the algorithmic turn in public administration is part of a longer legacy of quantification in government. In doing so, we echo earlier calls to situate the data-driven governance movement in government within the broader epistemology of quantification (Rieder and Simon, 2016).

So what then is at stake for the legitimacy of the administrative state? Administrative law has long emphasized the legitimacy and authority of agency experts by appealing to their efficiency, mitigating their capacity for arbitrary decision-making, and expanding their accountability (Bressman, 2003; DeCanio, 2021). We see echoes of such longstanding legitimacy practices in contemporary debates about agency uses of algorithms to administer policy.

Scholars have, for instance, emphasized algorithmic systems' capacity to enhance legitimacy through efficient and fair administration (Engstrom et al., 2020) and to make policy decisions visible (Levy et al., 2021; Coglianese and Lehr, 2019). More commonly, the literature has highlighted the challenges that algorithms pose to administrative legitimacy: making decisions more opaque (Pasquale, 2015), displacing decision-making power (Cooper et al., 2022b; Mulligan and Bamberger, 2019), threatening expert deliberation (Calo and Citron, 2020), and eroding public participation and oversight (Citron, 2007; Kroll et al., 2017).

That is, the literature on algorithms in public administration has primarily focused on *whether* algorithms advance or threaten legitimacy (Ajunwa, 2023; Calo and Citron, 2020; Waldman, 2019). In contrast with the broader quantification literature, little attention has been given to the particular circumstances under which algorithms are used to advance administrative legitimacy—and where they are not. We examine this question to illuminate *how* new technologies are not merely upholding or undermining existing static conceptions of legitimacy, but are also reshaping the enactment of legitimacy altogether. Despite the centrality (and tenuousness) of legitimacy, remarkably little empirical work has been done to examine how legitimacy is enacted in the administrative state.

3.1.1 Roadmap

We empirically investigate how the adoption of algorithms in public administration is transforming administrative legitimacy and its underlying values. Through a systematic qualitative analysis of reports and testimonies published by the U.S. Government Accountability Office (GAO), we examine how GAO conceptualizes the risks and promises of agency use of algorithmic technology and examine what values they point to in order to establish the legitimacy of incorporating algorithmic systems into governance processes.

To contextualize our analysis, first, we characterize the role of **legitimacy** in the administrative state. Then, we contextualize algorithmic technology use in U.S. public administration within the legacy of non-algorithmic forms of **bureaucratic quantification** and broader histories of administrative legitimacy.

We then introduce the **methods and empirical strategy**, through which we examine GAO's legitimacy practices as they relate to the domains, promises, and challenges of algorithmic use, and the interventions that GAO proposes. We organize our **findings** of how legitimacy is enacted and renegotiated around algorithms along the dimensions of *efficiency, non-arbitrariness,* and *accountability.* We close with a **discussion** of the ways that algorithms reshape and reorganize ambiguity in the administrative state, and the consequences for bureaucratic quantification–and for democracy.

3.2 Legitimacy Practices

We focus in particular on the U.S. administrative state's pursuit of legitimacy as a crucial site where legitimacy is negotiated through everyday practices. While in the U.S. the authority of legislators is legitimated by democratic elections, agency officials—who are appointed rather than elected—must look elsewhere. The history of U.S. agency policy-making points to the intricate strategies involved in establishing the legitimacy of the administrative state, centered around two key aspects of legitimacy: the suppression of arbitrariness and the expansion of accountability (Bressman, 2003). Attempts to minimize arbitrariness have relied on establishing objectivity through appeals to rules, procedure, and expertise as well as through demonstrably consistent and fair decisions. Meanwhile, efforts to expand accountability have involved transparency, public participation, and susceptibility to consequences. At times, the legitimacy of the administrative state has been conceptualized as a trade-off with the efficiency achieved through agency experts' flexibility and capacity to oversee complex issues (Bressman, 2003; Waldman, 2020). However, efficiency has alternatively been conceptualized as a unique affordance of the administrative state that justifies its existence and therefore plays a key role in its legitimacy (Calo and Citron, 2020; DeCanio, 2021). Examining the administrative state's legitimacy practices reveals the values for which legitimacy stands in: ***accountability***, ***non-arbitrariness***, and ***efficiency***. More importantly, looking closely at these legitimacy practices can reveal how the state interprets, assesses, and operationalizes these often contested values.

3.3 Administrative Legitimacy and Quantification

We propose that agency practices around algorithms belong to a class of legitimacy practices based in *quantification*. The "production and communication of numbers" (Espeland and Stevens, 2008, p. 402) has been taken up by STS scholars to understand how numbers are shaped and how they in turn shape the world. Understanding the adoption of quantification practices in administrative contexts has been central to the quantification literature (Mennicken and Espeland, 2019). In doing so, this literature examines the authority and production of authority of numbers (Porter, 1995; Deringer, 2018; Berman, 2022).

Notably, quantification is not only used to pursue legitimacy: it also changes how legitimacy is enacted. The proliferation of quantification in government has redefined legitimacy practices within policy spaces by privileging quantitative demonstrations of legitimacy. In the following subsections, we outline existing work about how quantification has been used to demonstrate legitimacy in bureaucratic governments. We highlight how quantitative practices have been used to pursue efficiency, non-arbitrariness, and accountability. We also discuss how quantification has changed the way that these values are conceptualized and enacted.

3.3.1 Quantification and Efficiency

Quantification has long been used toward the goal of more efficient government. Transforming policy decisions into numbers enables efficiency by simplifying complex policy landscapes, narrowing the language of dispute (Demortain, 2019) and enabling both the measurement of and optimization toward efficiency-based targets (Mennicken and Espeland, 2019).

Importantly, quantification allows policy-makers to measure, compare, and improve policies on the basis of their efficiency. Cost-benefit analysis and risk analysis are two forms of quantification that have flourished within U.S. agencies to demonstrate both the efficiency and legitimacy of policy decisions (Sunstein, 2002).

By simplifying complex information into easily comparable numbers, these quantitative comparisons enable dispersed surveillance and distant governance (Espeland and Stevens, 2008). This distance has been key to the efficient functioning of large-scale bureaucracies because it allows national governments to operate at scale, allowing the state to abstract away differences between individual subjects (1998a), excluding the complexities of local issues in service of efficiency (Demortain, 2019).

The close link between quantification and efficiency has transformed legitimacy practices to further emphasize efficiency over justifications of authority. For example, the administrative state has increasingly privileged efficiency-based analysis to demonstrate non-

arbitrariness (Porter, 1995; Sunstein, 2017). Similarly, Elizabeth Popp Berman (2022) describes the U.S. government's turn toward an economic style of thinking that privileges readily quantifiable and efficiency-based bases for legitimacy. This turn toward quantitative, efficiency-based legitimacy practices encourages economization, threatening to recast policy decisions as economic decisions and democratic goals as market goals.[2] Even when quantification fails to live up to its goals, these transformations are durable (Berman, 2022; Hacking, 1982) and have long-term implications for justifications of authority not based in efficiency.

3.3.2 Quantification and Non-Arbitrariness

Quantification promises to minimize arbitrariness through its suggestion of objectivity and neutrality (Espeland and Stevens, 1998a; Porter, 1995). Numbers derive authority from their perceived objectivity, resistance manipulation, and clarity— making them particularly useful in the political domain (Deringer, 2018).

The concept of "mechanical objectivity" is central to the way that numbers are used to demonstrate non-arbitrariness: suppressing subjectivity through rigid rules and procedures (Daston and Galison, 2007; Porter, 1995). This approach to non-arbitrariness aligns with efforts to constrain agency officials' discretion by legislating their behavior extensively. However, mechanical objectivity diverges from historical attempts to establish agency officials' legitimacy by appealing to their scientific expertise (what Theodore Porter (1995) calls "disciplinary objectivity"). Consequently, quantification has reshaped the meaning of arbitrariness within U.S. agencies: mathematical modeling of risk has taken the place of expert judgment as the gold standard for evidence against arbitrary decision-making (Porter, 1995).

Throughout the quantification literature, the objectivity of numbers is demonstrated through consistency and fairness. Consistency has been particularly essential for establishing the legitimacy of social numbers, demonstrating that these numbers are impervious to distortion by the individual whims of those who create them (Espeland and Stevens, 2008; Didier, 2020). Beyond consistency, quantification also promises to achieve fairness by eliminating human biases. As a result, quantitative demonstrations of unbiasedness have proliferated, in turn reshaping the meaning of fairness (Espeland and Stevens, 2008).

[2]On this Andrea Mennicken and Wendy Nelson Espeland (2019) write, "[N]umericization opens up not only new ways of political engagement; it also encourages economization by moving governments and citizens closer to principles of rational calculation, efficiency, competition, and market making, which in turn can run the risk of hollowing out democratic rule."

3.3.3 Quantification and Accountability

Contemporary statistical practices around open data and open science enact accountability first and foremost through transparency (Freese and Peterson, 2018). Under this paradigm, transparency enables verification, discourages incautious or dishonest practices through the threat of sanction, and allows claims to be compared and jointly assessed. Indeed, transparency has become closely associated with the production and communication of numbers that quantification enables, making visible decisions, people, objects that may previously have been hidden (Espeland and Stevens, 2008).

The transparency of numbers also facilitates accountability by enabling public participation in political life (Espeland and Vannebo, 2007). This participation is facilitated and through both political representation – via national census, for example– and debate. Indeed, William Deringer (2018) conceptualizes calculations as instruments of dispute, tools of argumentation and persuasion that have historically not only allowed those in power to exercise their authority, but have also allowed outsiders to critique those in power.

3.4 Methods

3.4.1 Study Setting

The U.S. Government Accountability Office (GAO) is an independent agency in the legislative branch that provides oversight over the federal government, including administrative agencies in the executive branch, on behalf of Congress. The GAO aims to preserve administrative legitimacy through its day-to-day operations, emphasizing nonpartisan information gathering and the accountability of the administrative state (GAO 2024b). As such, GAO offers a natural setting through which to explore everyday state legitimacy practices.

GAO pursues this mission via the publication of reports, either initiated externally by Congress (most common) or by existing laws, or internally via the Comptroller General. Congress has other oversight tools available–e.g., initiating committee hearings. Historically, Congress has leaned on GAO reports more heavily in times of unified government to advance shared policy goals, rather than to undermine the executive branch (Kennedy, 2023). These reports ultimately reflect areas of political attention, including specific ways that government functions become politicized (Figure 3.1).

Through their reports on agency use of algorithms, GAO implicitly intervenes on where and how it is legitimate for algorithms to take on administrative roles. These reports meaningfully shape agency practices: GAO estimates that the recommendations issued in their reports are taken up by agencies about 75% of the time. We use a corpus of reports published

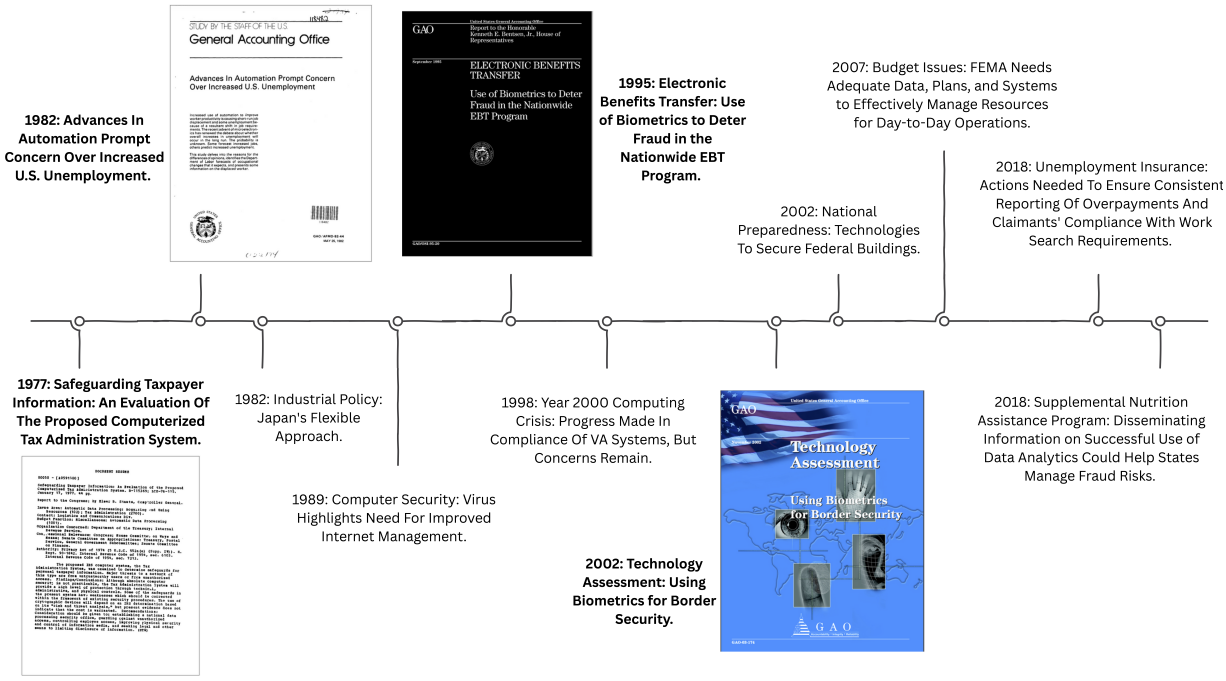


Figure 3.1: Illustrative examples of GAO reports in our sample. Political anxieties can direct attention to agency activities to target different concerns (e.g., fraud, unemployment, national security) and perform different values (e.g., efficiency, flexibility, consistency).

by GAO to examine everyday legitimacy practices and how they are shaped by, and shape, the algorithmic turn in the administrative state.

3.4.2 Data

Our analysis draws from a corpus of reports published by the U.S. Government Accountability Office since its inception. We selected all reports that included the term "algorithm" and/or the term "artificial intelligence." We chose these keywords to trace the evolution of contemporary algorithmic technologies that have reanimated questions of legitimacy in public administration. A broader search might also consider predecessors to these terms such as "automation" or "intelligent machines," but we limit our search for consistency in the types of technologies described. There were 953 reports that met our inclusion criteria, the earliest of which was published in 1968 (Figure 3.2). We randomly sampled 10 reports from each presidential administration in our corpus, except for administrations where fewer than 10 reports were published for which we included all reports (Table 3.1). This left us with a total of 90 reports, ranging in length from 11 to 888 pages. We performed qualitative coding on these 90 reports following the approach outlined in the following subsection.

Table 3.1: Number of reports included in final analysis by presidential administration. We randomly sampled 10 reports per presidential administration in our original corpus, except for the Johnson, Nixon, and Ford administrations in which fewer than 10 reports were published. For these administrations we included all available reports.

Administration	Number of Reports in Analysis
Johnson	1
Nixon	1
Ford	8
Carter	10
Reagan	10
George H.W. Bush	10
Clinton	10
George W. Bush	10
Obama	10
Trump	10
Biden	10

3.4.3 Empirical Strategy

We treat the GAO reports in our sample as *discourse materials* that both reflect and shape how legitimacy is and has been understood (see generally Clarke et al. (2017) on the analysis of discourse materials). To analyze these reports, we draw on abductive analysis (Timmermans and Tavory, 2022; Tavory and Timmermans, 2014) to identify themes around how GAO enacts legitimacy around algorithmic technology. By beginning from existing theories and looking for deviations from these theories, abductive analysis is designed to generate surprises and support the development of theory. We begin from Bressman's (2003) account of administrative legitimacy and the history of quantification and legitimacy. In our analysis, we contextualize our findings within the emerging algorithmic legitimacy literature, highlighting continuities and divergences from this literature.

We employed thematic line-by-line coding of each GAO report in our sample along the following dimensions:

1. Domain: Where does the government turn to algorithms? Where does GAO conceptualize algorithmic intervention as legitimate or illegitimate?

2. Promises: How does GAO justify the turn to algorithms? How is this transforming conceptions of legitimacy?

3. Challenges: Where does GAO conceptualize the risks and challenges of algorithms? How does this align with notions of legitimacy?

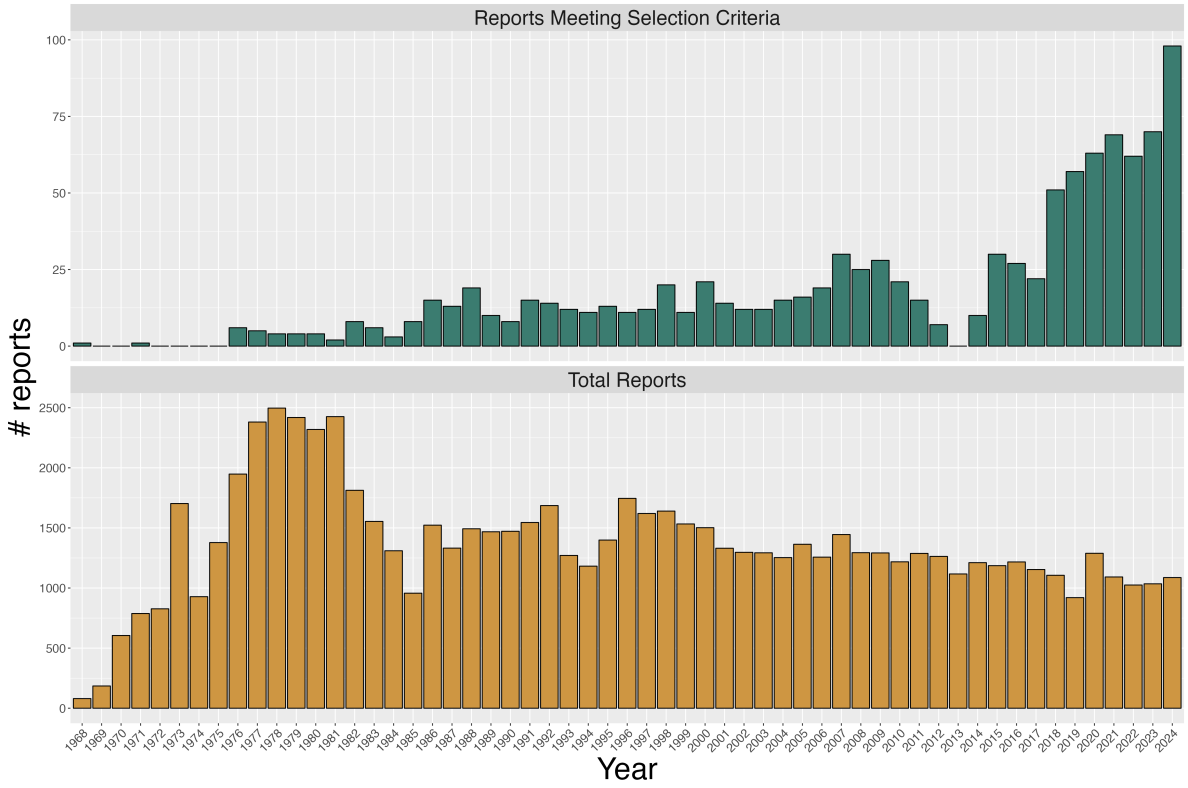


Figure 3.2: GAO reports 1968-2024. Top: Histogram of the number of GAO reports including the terms "algorithm" and/or "artificial intelligence" per year. (Note: our data collection stopped March 18, 2024, but this table includes all reports through the end of 2024.) Bottom: Histogram of all GAO reports during the same period for comparison.

4. Intervention: What steps does GAO propose to address the challenges outlined above? How are these interventions designed to ensure or demonstrate the legitimacy of algorithmic systems?

Applying this coding scheme, we might for instance find in a report "health care" as a domain, "speed" and "accuracy" as promises, "opacity" and "data privacy" as challenges, and "oversight" as a proposed intervention. Once we applied this coding scheme, we grouped related codes together and identified major themes. Specifically, we examine promises, challenges, and interventions through the lens of legitimacy, grouping codes under the major themes of accountability, efficiency, and non-arbitrariness – historical pillars of administrative legitimacy and quantification, which we find are also central to how GAO conceptualizes legitimate agency use of algorithms. Drawing on the history of administrative legitimacy, our analysis allows us to identify where algorithms extend, but also transform agency legitimacy practices vis-a-vis these three central dimensions. By examining how GAO discusses the promises and challenges of algorithmic systems, we highlight the aspects of legitimacy

that govern where and why algorithms are adopted. Attending to interventions designed to make algorithms more legitimate allows us to anticipate how conceptualizations of legitimacy are transformed by agency practices. We conceptualize these interventions as everyday legitimacy practices: the steps toward legitimacy that agencies take, not when legitimacy is in crisis, but rather when new modes of administration reconfigure the scope of legitimate administrative decision-making.

3.5 Findings

We analyzed how GAO reports between 1968 and 2024 have discussed agency use of algorithmic technologies and artificial intelligence. Through our coding and analysis of application **domains**, we found that the majority of agency uses of these technologies were adopted for internal management purposes– such as cybersecurity, billing, or agency workforce management– rather than for public-facing policy implementation. These findings match a 2023 GAO report showing that, except for "science," internal management was the most common agency use case of AI (U.S. Government Accountability Office, 2023).

Through our analysis of the promises and challenges of algorithmic technology we find that efficiency, non-arbitrariness, and accountability are still key values used to demonstrate legitimacy. However, we find that efficiency and non-arbitrariness are more central than accountability in establishing the legitimacy of algorithms. Finally, we examine the interventions used to address potential legitimacy challenges imposed by algorithm. We conceptualize these interventions as *legitimacy practices*, tactics used to enact legitimacy. Through our analysis of these interventions, we find that legitimacy is enacted in new ways, reflecting specific understandings of what makes agency algorithms legitimate.

In the following sections, we examine how the core values of administrative legitimacy are operationalized, emphasizing continuities and shifts with prior forms of quantification and the history of administrative legitimacy more broadly. To structure our findings, we draw on Bressman's (2003) account of legitimacy, which we divide into three dimensions: **efficiency**, **non-arbitrariness**, and **accountability**.

3.5.1 Findings: Efficiency in the Algorithmic Administrative State

In our sample of GAO reports, efficiency is frequently used to justify and legitimize the **promises of algorithmic systems** within agencies (Figure 3.3). Efficiency was the most commonly cited promise of algorithms, providing empirical evidence that efficiency narratives

Themes & Sample Quotes About **Efficiency** as a Value		
Promises	Efficiency in economic terms, including as a value that can be expressed quantitatively	"The Redistribution and Disposal System (RADS) is a computerized mathematical model designed to serve as a management tool for decisions concerning the reallocation of overstocked items to other GSA [U.S. General Services Administration] depots or to disposal. The heart of the system is an algorithm that determines the **economically optimum** set of allocations throughout the GSA system for all the stock of each item being analyzed." – *Federal Information Sources & Systems,* 1976 "Compliance programs will **generate revenue** by using market segment analysis to select the **most productive cases** and by providing frontline enforcement personnel with immediate access to taxpayer data files and automated tools to perform both complex calculations and routine clerical tasks **more quickly and accurately.** This will result in **lower costs to both the taxpayer and the Service**." – *Financial Audit: Examination Of IRS' Fiscal Year 1994*, 1995 "ACI [the Automated Criminal Investigation System] will permit Criminal Investigation employees to **more efficiently** and effectively complete investigations, using information obtained from sources throughout the law enforcement community. It is projected that the automation of work practices through ACI will **increase employee productivity by as much as 20%**." – *Financial Audit: Examination Of IRS' Fiscal Year 1994,* 1995
	Efficiency as functionality (increased performance, accuracy)	"Being able to predict taxpayer behavior using statistical models helps improve the likelihood of a **desired outcome** by matching the outreach to a given taxpayer's characteristics. The IRS is better able to **efficiently use limited resources to accomplish tax administration goals** such as fairness, reducing taxpayer burden, revenue collection, and ensuring compliance through the judicious use of these data analytic techniques." – *Financial Audit: IRS's FY 2019 & FY 2018 Financial Statements,* 2019
	Efficiency as necessity due to e.g., size, complexity	Computer system support is **essential** to collect, analyze, and transmit the **vast array of information** used to plan and manage a service's enlisted force." – "*Spaces," "Faces," and "Bucks": An Overview of the Role of Automated Information Systems in Managing the Enlisted Force,* 2010 "IRS converted to automated data processing (ADP) because statistics showed that the Service's workload was increasing **beyond the capacity** of conventional manual and machine processing capabilities." –*Safeguarding Taxpayer Information: An Evaluation Of The Proposed Computerized Tax Administration System,* 1977
Challenges	Efficiency undermined by economic costs of algorithms	"All too often, the lure of the electronic office has led agencies to launch **expensive automation programs** without adequately defining what characteristics the automated system should have, which alternative approaches for achieving automation are most efficient and effective, or even how best to use computer and telecommunications equipment once it is acquired." –*The GAO Journal, No. 1, Spring 1988,* 1988
Interventions	Non-algorithmic quantification	"A detailed **cost-benefit analysis** must be conducted to determine that the benefits gained from a system outweigh the costs." –*Aviation Security: Challenges in Using Biometric Technologies,* 2004

Figure 3.3: **Themes of efficiency** as a value emphasize the promises of algorithmic adoption across a range of expressions of efficiency, but also appear as a potential challenge of algorithmic adoption and in interventions for establishing legitimacy through other means.

have been core to the uptake of automated systems. This finding aligns with legal and sociological scholarship that theorizes that algorithmic decision-making systems' orientation toward agility, frictionlessness, and cost-reduction make them a natural ally of the U.S. government (Waldman, 2020)—i.e., of a government that increasingly prioritizes efficiency over other social and political values (Berman, 2022). However, a closer look at this rhetoric in the GAO data reveals the *malleability* of efficiency: the many ways it can be evaluated, and the many values that it stands in for. The promises of efficiency are expressed in different ways: as economic, as functionality, and as necessity. When efficiency is used to legitimate the adoption of algorithmic technology, we observe a tension between cost-reduction and functionality (getting the job done cheaply; getting the job done at all).

The reports primarily conceptualize efficiency in terms of *economic* optimality, discussing

efficient allocation, revenue generation, cost reduction, and enhanced worker productivity. In government contexts, scholars have argued that by reducing the need for human deliberation, automated decision-making systems promise to save money by reducing labor costs, and to save time by shortening the decision-making process, freeing up human decision-makers to apply their efforts to other tasks, and ultimately reducing backlog at over-burdened federal agencies (Engstrom et al., 2020). This matches the arguments used by the GAO, e.g., to "efficiently use limited resources" to accomplish goals towards their "desired outcome" by adopting algorithms (GAO 2019, Figure 3.3). In addition to matching existing arguments about government algorithms,this emphasis on economic measures of efficiency also echoes how other forms quantification has been used to legitimate the administrative state. Notably, the proliferation of cost-benefit analysis and economic professionalization in agencies have contributed to a vision of the administrative that is justified first and foremost by its cost-effectiveness (Porter, 1995; Berman, 2022).

Despite its prevalence, economic cost is not the only conception of efficiency used to justify the promises of algorithmic technology. Beyond reduced costs and increased revenue, GAO reports also emphasize the promise of enhanced performance and needed functionalities, pointing to the effectiveness and accuracy of these new systems. However, this functionality remains secondary to economic efficiency improvements. GAO sometimes presents the efficiency of algorithmic systems as **necessary** due to the size or the complexity of agency problems (Figure 3.3). In these circumstances, GAO relies on efficiency alone to establish the legitimacy of adopting algorithmic technologies. Where algorithms are framed as necessary to meet the demands of increased scale and speed, efficiency risks crowding out other legitimacy considerations, like functionality. Prior work about the adoption of algorithmic systems in government, prioritizing economic efficiency over other values risks undermining the administrative state's pursuit of legitimacy, particularly where the adoption of algorithmic systems implicates human rights (Waldman, 2020). While this literature largely discusses how algorithms *undermine* legitimacy, we argue that instead legitimacy is being *reinterpreted* to emphasize economic efficiency over other justifications of agency authority.

Similarly to other forms of quantification, enhanced efficiency is the central promise of algorithmic adoption (Engstrom et al., 2020). Here, we observe continuity in the way that efficiency is operationalized: algorithms promise to extend the legacy of quantification by automating the assessment of and optimization toward economic goals. Yet, GAO also engages with efficiency in how they **conceptualize challenges of algorithms** through the economic costs of algorithms. We observe that the promise of greater efficiency is uncertain. When presented as a risk of algorithms, efficiency continues to be conceptualized in economic terms. Our findings indicate that, while efficiency is conceptualized in similar ways under the

algorithmic regime and the pre-algorithmic quantification regime, there is less certainty about the result of algorithmic technology. Although economic efficiency remains an essential tenet of legitimacy, algorithms promise to further efficiency just as they threaten to undermine it.

To **intervene** where efficiency is uncertain, we find that GAO relies on non-algorithmic forms of quantification to establish the legitimacy of new forms of algorithmic administration. To mitigate the efficiency risks brought on by the uncertainty associated with algorithms, GAO turns to more established everyday quantification practices, like cost-benefit analysis or quantitative risk analysis, as demonstrated in Figure 3.3. We note the collision of long-time legitimacy mechanisms, like cost-benefit analysis and oversight, with new mechanisms like technical validation, pointing to a subtle shift to extend rather than disrupt prior logics of cost-based quantification. By combining the adoption of algorithms with more established forms of quantification, GAO attempts to reinforce algorithmic technology where uncertainty around its efficiency threatens to undermine administrative legitimacy. This reflects continuity in conceptions of efficiency and legitimacy between algorithmic and non-algorithmic forms of quantification. However, algorithmic technologies are not associated with the same level of legitimacy as other forms of quantification like cost-benefit analysis, which has been institutionalized within administrative legitimacy practices.

3.5.2 Findings: Non-Arbitrariness in the Algorithmic Administrative State

GAO refers to the value of non-arbitrariness as a **promise of algorithms** (Figure 3.4). As with other forms of bureaucratic quantification, the value of non-arbitrariness is demonstrated through appeals to *consistency* and *objectivity*. Moreover, we observe that this non-arbitrariness afforded by algorithms is positioned in contrast with expert judgment.

The administrative state's efforts to establish its own legitimacy have long appealed to their own objectivity and non-arbitrariness, including quantification as an instrument for objectivity. However, the critical literature stresses that—just as any decision-making system—algorithms are never neutral. Rather, algorithms require policy choices around "what data sets to use to train the algorithm, to how to define what the algorithm's target output is, to how likely the algorithm is to produce false positives versus false negatives" (Kaminski, 2020, p. 122). That is, algorithmic technology cannot and can never achieve objectivity, although their proposed benefits include reduced bias, increased accuracy, and increased consistency (Engstrom et al., 2020; Kleinberg et al., 2018, e.g.). We see these values – fairness, accuracy, and consistency—standing in for evidence of non-arbitrariness in algorithmic decision-making.

Themes & Sample Quotes About **Non-Arbitrariness** as a Value		
Promises	Consistency (typical of other types of quantification)	"For example, several models are used primarily for forecasting a single crop, a series of crops, or an entire agriculture sector over a period of time. Used to assess probabilities of future outcomes, such models can provide policymakers with **systematic** speculations about the future." *–Food And Agriculture Models For Policy Analysis,* 1977
	Objectivity (in contrast to expert judgment)	"[Ground officer positions] are not based on measured work load but on what the Navy refers to as **a ground officer 'algorithm.'** While **this term suggests a systematic mathematical approach** to developing requirements, as used here it represents a list of positions the Navy believes are required, **based largely on corporate management judgment**. … A 1981 [Chief of Naval Operations]-contracted study of the Navy's development of officer requirements criticized the Navy's methodology, stating that the 'methodology used is **not sufficiently quantitative in nature** to **provide adequate justification** for the ground officer requirements.'" *–Navy Manpower: Squadron Manpower Program Needs Improvement*, 1987
Challenges	Risks of inaccuracy (legitimacy acquired through accuracy)	"Although simulations are useful tools, they rely on data that may be incomplete and **assumptions that may be inaccurate**. Constructing an accurate simulation requires that the behavior of what is being simulated be well understood; there is a great danger in accepting the results of computer simulations as representing reality, rather than using them as design tools. **The accuracy of a simulation can only be checked by comparisons with measured results in the real world, a process called 'validation.'**" *–Strategic Defense Initiative: Estimates Of Brilliant Pebbles' Effectiveness Are Based On Many Unproven Assumptions*, 1992
Interventions	Earlier reports focus on humans in the loop (i.e., expert judgment) to mitigate algorithmic arbitrariness	"Finally, systems analysis is **not a substitute for the decisionmaker's judgement**. It is an aid--a tool--that can be used by the decisionmaker to sharpen his judgment." *–Glossary for Systems Analysis*, 1968 [describing an algorithmic driving simulator:] "(1) the motion envelope for any simulator will impose inherent limitations on some applications because **there is no analytical, scientific, or engineering way to define the required motion envelope for any given application** and (2) the size of the motion envelope will still **rely on simulation experts' best guess**, even after washout algorithms have been used." *–Motor Vehicle Safety: Key Issues Confronting The National Advanced Driving Simulator*, 1992
	Later reports focus on humans as reviewers of correctness, not substance, of algorithmic outputs	``[W]e used **professional judgment** to determine whether cases with typos were sufficiently similar to consider as matches" *–2014 Lobbying Disclosure: Observations On Lobbyists' Compliance With Disclosure Requirements*, 2015 "The first step in the **manual review process** is a triage review, which is intended to identify loans of less than $2 million that can be easily resolved by determining that an automated screening alert was **invalid**, according to loan review contractor officials." *–Paycheck Protection Program: SBA Added Program Safeguards, but Additional Actions Are Needed*, 2021

Figure 3.4: **Themes of non-arbitrariness** are called upon to address the promises, challenges, and legitimacy practices around the adoption of algorithmic technologies.

Consider one example, where GAO takes issue with a system referred to as a the "ground officer algorithm" for its misleading name. Unlike mathematical or computerized algorithms, which GAO claims are "systematic," the ground officer algorithm relies primarily on human judgment, which they imply does not entail the same protections against arbitrary decisions (Figure 3.4). In addition, we see that quantitative demonstrations of non-arbitrariness are privileged over judgment; for instance, quantitative information takes the place of, and makes obsolete, other forms of justification. The adoption of algorithms extends quantification's turn toward mechanical objectivity as a way to enact non-arbitrariness (Daston and Galison, 2007), further marginalizing the role of agency expertise in demonstrations of administrative legitimacy.

While non-arbitrariness is sometimes framed as a promise of algorithms, it also appears as a **challenge of algorithms in the administrative state**, focused particularly on potential inaccuracies. We find that where the arbitrariness of algorithmic systems is in question, demonstrations of non-arbitrariness are based in appeal to metrics like accuracy rather than reason-giving. Providing expert-based justification and reasoning for deviating from prior decisions has long been a cornerstone of administrative defense against the arbitrary and capricious legal standard to which agencies are held. However, in the context of algorithms, accuracy metrics stand in for more reason-based explanations for decisions. Under this vision of non-arbitrariness, it is not the unjustified exercise of discretion, but the possibility of errors that threatens to undermine the legitimacy of algorithmic systems. In one example, we observe non-arbitrariness conceptualized as predictive validity, refocusing attention to prediction and accuracy metrics.

This is distinct from other quantitative demonstrations of non-arbitrariness like cost-benefit analysis: Cost-benefit analysis leaves room for agencies to engage in interpretation and judgment through deciding how to assign value, an inherently ambiguous task. Yet algorithms inherit the same ambiguities, particularly in deciding how to operationalize and encode policy goals. These acts of interpretation become detached from assessments of legitimacy when evaluation becomes about assessing accuracy, a task which assumes away ambiguities in favor of a simple binary between correct and incorrect.

Focusing on inaccuracy as the primary risk of non-arbitrariness represents a shift in how arbitrariness is understood within agencies. Even when algorithmic systems are highly accurate, the inner workings of systems may be opaque or unexplainable in ways that challenge traditional notions of non-arbitrariness. As Emily Berman explains: "To lack arbitrariness, the government must have some justification for the exercise of its power, and that exercise of power must have some connection to the proffered explanation" (Berman, 2018). Linking non-arbitrariness with reason-giving reveals *why* authority is exercised– whether this author-

ity is in line with existing law, in line with precedent, or in line with the public good. However, these imperatives of reasoned decision-making are thwarted when administrative algorithms lack transparency and explainability. In sum, algorithmic non-arbitrariness demonstrations focus visibility on system accuracy as opposed to the assumptions and decisions embedded within system design.

This shift in non-arbitrariness is further exemplified in GAO's legitimacy practices. Across the GAO reports, we observe that **interventions** to identify algorithmic errors emphasize human assessment, but that the role of human judgment shifts over time. In our earliest report—written in 1968 and quoted in Figure 3.4, GAO privileges human judgment above algorithmic outputs. This orientation toward human judgment continued for decades. In this context, human expertise, i.e., expert judgment, is positioned as a solution to the problem of inherent uncertainty, ambiguity, or general purpose-ness that potentially arbitrary machine decisions leave open.

However, in more recent reports, we observe technical improvements become a primary intervention while human judgment is reduced to the role of review *after* an algorithm has already made a determination. In this case, judgment is only used to assess whether an algorithm is "correct," while broader conceptual ambiguity is unacknowledged. Human judgment is brought in to rectify algorithmic oversights in situations where there is an assumed "correct" answer, rather to provide a more nuanced assessment (note for instance, the change from "expert" judgment in 1992 to "professional" judgment in 2015). Thus, even in situations where human intervention is proposed to legitimate algorithms, the expert status of agency officials is minimized in light of a more rote, objective set of duties vis-à-vis algorithmic oversight. As a result, the role of judgment is diminished under an emerging set of legitimacy practices that conflate non-arbitrariness with accuracy.

3.5.3 Findings: Accountability in the Algorithmic Administrative State

We find that accountability is not called upon to legitimate the adoption of algorithmic technology to the same extent as efficiency and non-arbitrariness. In the GAO reports, we do not observe accountability called upon to support the adoption of algorithmic technologies; instead, accountability appears primarily in the context of their **challenges.** The examples in Figure 3.5 demonstrate the challenges of oversight due to both opacity and inherent technical limitations.

Critiques of algorithmic systems have raised concerns that the inscrutability of complex technical systems poses a threat to transparency, explainability, and, as a result, on the

Themes & Sample Quotes About **Accountability** as a Value		
Challenges	Algorithms seen as imposing challenges for accountability	"This **lack of transparency** limits the ability of auditors and others to detect error or misuse and ensure equitable treatment of people affected by AI systems." – *Artificial Intelligence: An Accountability Framework for Federal Agencies & Other Entities*, 2021 "Although some scientists and engineers believe that the needed [algorithms and software] for this purpose will be available in the near future, others believe that the software required is unattainable and that **such a system cannot be trusted because it can never be adequately tested**" – *Strategic Defense Initiative Program: Better Management Direction And Controls Needed*, 1987
Interventions	Emphasis on documentation (similar to algorithmic accountability's emphasis on transparency artifacts)	"**Documentation of information systems** is critically important for solving other information systems problems. For example, systems cannot be integrated or personnel trained if it is not clear exactly what a specific system does and how it can be operated."– *"Spaces," "Faces," and "Bucks": An Overview of the Role of Automated Information Systems in Managing the Enlisted Force*, 2010 "Entities should **document requirements—including performance metrics**—for the Al system throughout the life cycle. Entities should **document methods to assess performance**—which can include input-output tests, stress tests, and evaluations of model drift— to ensure AI systems meet their intended goals. Entities should **provide access to performance test results, change logs, and other documentation describing updates and key design choices**, and **provide a copy of the model or algorithm code to third-party assessors** of Al systems." – *Artificial Intelligence: An Accountability Framework for Federal Agencies & Other Entities,* 2021

Figure 3.5: **Themes of accountability** in the GAO reports emphasize the challenges of algorithmic adoption and how legitimacy practices could achieve accountability via documentation.

ability of non-technical stakeholders and experts to meaningfully participate in accountability processes. The lack of explicit attention to accountability from the GAO may be related to the fundamental barriers to accountability around algorithmic contexts. Building upon Helen Nissenbaum's (1996) "Accountability in a Computerized Society," A.F. Cooper and colleagues (2022b) have theorized that the barriers to accountability are heightened in an increasingly algorithmic state. In response to accountability concerns, a significant portion of the literature on the use of algorithms in government has focused on how these systems can be made more accountable. These solutions tend to focus on either transparency and explainability or on expert and public participation. (Citron, 2007; Vaccaro et al., 2019; Kaminski, 2020; Busuioc, 2021; Ajunwa, 2023)

Some research argues that algorithmic systems, especially those designed to be explainable, enable or even improve transparency in government decision-making (Coglianese and Lehr, 2019; Kleinberg et al., 2018; Levy et al., 2021) through explicit problem formulation and explanation of design choices. In contrast, we observe that accountability-related **promises** are seldom used to justify the adoption of algorithms; accountability is frequently cited as a **challenge** of algorithmic adoption. And while opacity and transparency concerns are highlighted by GAO, it is worth noting the well-recognized risks to algorithmic technologies that were not highlighted in GAO reports. Work on data justice highlights the expansion of data-driven surveillance that undergirds algorithmic governance (Dencik

et al., 2016; Kitchin, 2023). Prioritizing efficiency while neglecting accountability legitimizes algorithmic public administration, without imposing a duty to justify the intensified data collection that these algorithms demand.

This focus on opacity as the primary accountability challenge is reflected in GAO's practices around algorithms. We observe that **interventions** and legitimacy practices based on accountability emphasize documentation. This emphasis aligns with the algorithmic accountability literature's embrace of transparency artifacts. Indeed, treating documentation as a substitute for accountability can be seen as part of a larger policy shift towards traceability as accountability (Kroll, 2021, e.g.). This shift seems to be an interesting outcome of algorithmic adoption: This new conception of accountability is a shift from traditional administrative legitimacy practices that have centered presidential or judicial oversight and public participation as key means to establish accountability (Bressman, 2003). While documentation might in theory facilitate such oversight and participation, this link is not made explicit in the GAO reports that we analyze. Instead, "third-party assessors" are invoked Figure 3.5, perhaps reflecting the perceived challenges of ensuring algorithmic accountability without technical expertise. Archival scholars have long noted the importance of open records for democratic accountability, but also highlight the complications brought about by questions of control, access, trustworthiness, and preservation (Cox and Wallace, 2002). These concerns are particularly salient when information is governed by private entities, such as algorithm developers, with different obligations to the public than traditional government actors (Kranich, 2006; Baykurt, 2025). While GAO has identified transparency as a first step toward accountability, they risk overlooking these challenges when documentation practices are not explicitly designed with goals like democratic participation in mind.

3.6 Discussion

Quantification has long been used as a strategy to produce legitimacy. In turn, quantification has shifted how legitimacy is enacted by the state. Our findings show that the the uptake of algorithms is an extension, not fully a deviation from, this history of quantification. Quantification's promises of efficiency, non-arbitrariness, and accountability are similarly used to legitimate agency adoption of algorithms and artificial intelligence. We observe that the promises of algorithms follow similar arguments as quantification—but are also tempered by accompanying concerns about how they might instead undermine efficiency, non-arbitrariness, and accountability.

Unlike prior forms of quantification that have been adopted to produce legitimacy in contentious, low-trust settings (Porter, 1995), our findings show that algorithms have instead

been adopted for uses that are not subject to significant external scrutiny. We find that the majority of domains for which algorithmic adoption was discussed were related to internal agency management, in line with reported agency use of AI (U.S. Government Accountability Office, 2023).

This key difference—algorithms appearing in internal use rather whereas quantification historically intervened upon contested, low-trust settings—points to a subtle deviation in the administrative use of algorithms. We propose that algorithms have not yet been institutionalized as legitimation tools in U.S. public administration. For the administrative state, legitimacy is both used as a reason to adopt algorithms, and as a downstream impact of algorithms.

Through this dual function, administrative algorithms shift what makes bureaucratic authority legitimate: the adoption of algorithms encourages legitimacy practices that align with the affordances of these technologies. Like earlier forms of quantification, algorithms encourage number-based, easily measured demonstrations of legitimacy, like financial cost. Consequently, algorithms are positioned as a solution to problems of efficiency—particularly of scale and cost— rendering increased efficiency the primary outcome upon which the success, and the legitimacy, of algorithms is assessed. Algorithms also bring about new assessments of legitimacy like the use of predictive accuracy, a traditional way of evaluating algorithmic technologies, to stand in for other demonstrations of non-arbitrariness. Simultaneously, algorithms have a downstream impact of making accountability more challenging. Where algorithmic technologies don't align with traditional legitimacy practices—notably, their opacity contrasts with traditional demands for transparency, participation, expert reasoning, and accountability—these concerns are often sidelined. This encourages algorithmic adoption in less visible arenas, and ultimately sidelines accountability as a central tenet of legitimacy within the algorithmic administrative state.

Beyond shaping legitimacy based on what is actionable, these technologies also rearrange organizational function. Quantification not only shapes what practices stand in for legitimacy, but also how they are negotiated. Algorithms do the same. Scholars and government agencies have called to attend to not just what algorithms do, but how they claim to do it. Specifically, the implicit measurement processes in quantified and algorithmic systems change where administrative decisions are made: What are these systems aiming to do? What concepts (benefits eligibility, fraud; predicted relevance; predicted needs) are operationalized? How, and by whom? (See, e.g., Jacobs and Mulligan (2022).) As Deirdre Mulligan and Kenneth Bamberger (2019) have argued, procurement of algorithmic technologies can swiftly displace and redefine administrative functions (and expertise, accountability, and oversight). This literature reveals that replacing one component of a sociotechnical system

(for example, whether an agency official or an algorithmic system performs the task of policy administration), even when ostensibly performing the same function, can have larger political implications for the system's values (Mulligan and Nissenbaum, 2020; Goldenfein et al., 2020; Abdu et al., 2024). Understanding how algorithmic technologies reshape governance processes remains an open empirical question (Farrell, 2025).

Both existing administrative practices and the technical affordances of algorithmic technologies come together to shape how algorithms are incorporated in government. The resulting collision has meaningfully shaped how agency legitimacy is understood; algorithmic technologies are merely one part of a growing push toward efficiency. Although efficiency comprises many meanings, a cost-based conception of efficiency has served as one of the strongest rhetorical strategies for both deregulatory activity and algorithmic adoption. Efficiency, including lack of so-called waste, is used to demonstrate the legitimacy of agencies. However, this perspective—which evaluates agencies not by the tasks they accomplish but by how little they cost—invites systematic underfunding of the administrative state (itself not a new activity)[3] and its dissolution (words containing "efficient" or "efficiency" appear 154 times in Project 2025 (Dans and Groves, 2023); cf. Chen (2025).)

This efficiency imperative acts as a double-edged sword with which the administrative state can never establish its own legitimacy. The rise and fall of the U.S. Digital Service (USDS) offers an emblematic cautionary tale. Founded in the Obama era, their mission was to "deliver better government services to the American people through technology and design," in part by "continuously improv[ing] critical government services... Partner[ing] with federal agencies to scale digital best practices" (Digital Service, accessed April 21, 2025). The attempted rebrand of USDS under the Department of Government Efficiency (DOGE) moniker might seem logical, particularly as a move towards an "AI-first strategy" that extends this long-standing move toward quantification, it led to the dismantling of the USDS and the overall contraction of administrative agencies (Kelly, 2025). That is, the same tools used to advance toward administrative legitimacy can be used to attack it (Chen, 2025). As Henry Farrell and Abraham Newman (2025) recently wrote, Elon Musk "is turning the U.S. government's own strategy against itself." Scholars, agency officials, and citizens concerned with the adoption of algorithms and AI in public administration must reckon with the broader movement toward efficiency-based legitimacy practices that algorithms accelerate and entrench.

[3]Agencies like GAO tend to only get attention when they are being weaponized. Ronald Reagan infamously declared the urgency around his appointment for Comptroller General of the GAO "to restore the faith of the people in the integrity and efficiency of government. This is one of the most important appointments that I shall make as President" (Reagan, 1981).

3.6.1 Limitations

We acknowledge some basic limitations to our empirical strategy. First, our analysis is necessarily limited by its U.S-centric approach, and within that context, GAO is only one aspect of how legitimacy is operationalized. We could instead consider other empirical sites for understanding legitimacy in the administrative state, e.g., through other oversight bodies, in judicial decisions, or across individual agencies and their on-the-ground deployment. Choices comprising our empirical strategy, including case studies, qualitative coding, limited sample (random subsample of identified corpus), and a non-representative sample due to choice of corpus (reports with "algorithm" or "artificial intelligence"), all limit our potential findings. Although we limited our search to only a narrow set of keywords, our findings cover a wide range of technologies. Early technologies include biometric identification, defense simulations, and cryptographic cybersecurity, while more recent reports cover risk prediction algorithms and automated compliance tools. The functionality and ethical implications of these technologies varies widely. Despite this variation, we observe widely similar legitimacy language across different domains and types of technology. Additional text-based methods including larger-scale computational analysis could additionally be used to expand this analysis.

Yet, GAO's central role in overseeing the entire administrative apparatus of the U.S. government provides an important case study in how everyday practices around technology shape what is seen as legitimate. Through their oversight and policy recommendations, GAO governs when it is permissible for agencies to turn to algorithms, on what grounds, and with what safeguards. More broadly, the institutional context of GAO, the historical turn to quantification in public administration, and the technical capabilities of algorithms all affect what is ultimately seen as a legitimate agency use of algorithms. This results in a co-constitutive relationship between administrative legitimacy and agency adoption of algorithms.

Conversely, existing legitimacy practices shape how algorithmic technologies are incorporated into agency workflows. GAO's long history as an accounting institution, for example, has shaped the way that it performs its oversight duties. Even as its mission has shifted to emphasize accountability, GAO conceptualizes accountability's primary goal as ensuring the cost efficiency of agencies.[4] Already institutionalized everyday legitimacy practices like cost-benefit analysis reflect a broader push toward administrative efficiency and dictate where agency algorithms are seen as legitimate. When algorithms are financially costly, they are

[4]The GAO (2025) self-describes: "GAO examines how taxpayer dollars are spent and provides Congress and federal agencies with objective, non-partisan, fact-based information to help the government save money and work more efficiently."

seen as risky, but when they help under-resourced agencies increase productivity they are seen as necessary despite other risks and limitations. This accelerates an ongoing collision of corporate and administrative legitimacy that evaluates agencies using the same metrics as private companies, overlooking the government's primary goal of serving the public interest (Short, 2023). We see the role of agency officials shifting from experts to professionals with technical expertise increasingly expected to be located *outside* of agencies, opening the door even wider for corporate participation in policy-making and inviting further study into non-governmental legitimacy practices.

3.7 Conclusion

We contextualize the adoption of algorithms in U.S. public administration within the legacy of non-algorithmic forms of quantification and broader histories of administrative legitimacy. Whereas the history of bureaucratic quantification points toward the deployment of numbers to confer legitimacy, our analysis shows that government algorithms have largely not been deployed to advance administrative legitimacy. Yet, the justifications used to deploy algorithms in agencies' everyday activities demonstrates how legitimacy is reformulated to account for new practices. We show that while the traditional pillars of administrative legitimacy—accountability, non-arbitrariness, and efficiency—remain essential, their relative importance has shifted alongside the practices associated with them. Specifically, we note that algorithms are legitimized by diverting attention away from administrative accountability practices and refocusing legitimacy conversations to highlight the importance of efficiency.

Changing notions of non-arbitrariness that center accuracy over explanation transforms the role of experts from substantive, conceptual work to supervisory auditing.

While we have focused on the administrative state as a key site where legitimacy, authority, and power are negotiated, our findings have implications beyond the state. Platform studies scholarship has noted that private technology companies have increasingly taken on government functions, extending questions of legitimacy beyond the state and further complicating questions of accountability (Iliadis and Acker, 2024; Cohen, 2025). As authority is increasingly wielded by those outside of government, the legitimacy of their authority, and the practices used to establish their legitimacy, must be continually re-assessed. In this paper, we demonstrate the need to examine legitimacy not only as a theoretical construct, but as an empirical boundary defined and maintained through the everyday practices of those in power. This has important implications for who has the authority to enact policy, and on what grounds—both within the administrative state and beyond.

CHAPTER 4

Ethical Closure Problems in Formalizing Technologies[1]

4.1 Introduction: "Perfect Technologies of Justice"

Administrative officials across U.S. government agencies serve the important but complex function of implementing policy. Despite its important role in operationalizing policy "on the ground" (Edelman, 2020), the administrative state is plagued by frequent challenges to its legitimacy. Unlike publicly elected representatives, administrative officials are appointed: their role as policymakers (through regulation and policy implementation) is not legitimated through democratic elections. Instead, the state has historically sought to justify agency authority by other means, emphasizing non-arbitrary interpretations of policy, accountable decision-making, and efficient operations (Bressman, 2003). Administrators' implementation decisions are at times desired to align with a policy's intended goals and values; with citizen expectations and political machinations; and with institutional expectations, all while respecting both democratic and technocratic norms of agency policy-making (Bamberger, 2008, 2009). When such values, expectations, and norms conflict, agencies engage in legitimacy-making exercises to justify their continued policy-implementing capabilities. In this paper, we investigate how mathematical and algorithmic formalisms—specifically what we call formalizing technologies—intersect with such conflicts. We show how such technologies are used to resolve these conflicts to perform legitimacy, but also reshape and create conflicts anew.

We introduce the term *formalizing technologies* (FTs) to refer to mathematical frameworks intended to operationalize and enact policy. We define formalizing technologies as mathematical frameworks that implement policy ideals, such as efficient markets (e.g., through mechanism design), fair allocation, certain security goals (e.g., through cryptogra-

[1]This chapter is drawn from the following paper, which is under revision at the *ACM Journal on Responsible Computing*: **Amina A. Abdu**, Jeremy Seeman, and Abigail Z. Jacobs. "Ethical Closure Problems in Formalizing Technologies."

phy), or democratic participation (e.g., through voting). Drawing on existing uses of formalizing technologies in government, we build upon the science and technology studies and quantification literature to explore the mechanisms through which formalizing technologies create legitimacy—and the consequences of this legitimacy for policy-making. Specifically, we show how these formalizing technologies enact *closure*: closure of social problems, of policy solutions, of values and of ethics in the political context.

Here we build on multiple strands of literature. First, we show how multiple theories of *scientific closure* in organizational theory and science and technology studies speak to the ways that consensus and legitimacy are formed around comparable technologies. Because we explore technologies that are used to shape social, policy-oriented outcomes, we also look to how technologies in the administrative state produce closure around values. In a context where legitimacy is at stake and a lack of consensus can undermine agency authority, it is important to understand where this legitimacy comes from, and where it is sensitive to challenges. We connect closure to the production and performance of legitimacy, whereby consensus provides moral justification for the exercise of power.

Shifting to the context of the administrative state, liberal political theory presupposes legitimacy as a necessary condition for just functioning of the administrative state (Rawls, 2020a). Seminal works in this literature identify three success criteria for justifying productive political dialogue at the heart of U.S. administration: first, rationality, the ability to explain a resource allocation decision; second, consistency, the consistent application of rules across parties and circumstances; and third, neutrality, the equal moral standing of different conceptions of the good and their advocates (Rawls, 2020b). In combination, these three criteria can be operationalized by what Bruce Ackerman (1980) calls "perfect technologies of justice", imagined utopian means to implement policies without obstacles, distortions or costs. Such technologies dictate the terms and processes of democratic deliberation, not the content of the deliberation itself, as the point of connection between democratic legitimacy and just administration (Peter, 2007). Doing so inextricably links these success criteria to the efficient functioning of the state: in Ackerman's view, an efficient liberal state is necessarily a just state.

We argue that formalizing technology (FT)s most closely align with this liberal utopian vision by scaffolding the terms of democratic deliberation in a technically productive manner. The use and rhetoric of FTs aligns with liberalism's governance goals of accountability, efficiency, and non-arbitrariness. We explore the rhetorical machinery through which formalizing technologies produce ethical closure, and its implications for policy. (Even the language of *technologies*, rather than policy instruments, serves to distance the tool from its political and social function.)

The paper proceeds as follows. In Section 4.2, we introduce and define *formalizing technologies* (FTs), providing examples and discussing the affordances that make FTs useful in the pursuit of legitimacy. In Section 4.3, we synthesize existing frameworks of scientific closure, focusing on building consensus and preserving autonomy in science. We subsequently examine how an analogous framework of *ethical closure* describes the process by which policy implementations impose consensus and autonomy in decisions about essentially contested values. In Section 4.4, we examine how the scientific and ethical closure process plays out through FTs, focusing on two case studies: formal privacy enhancing technology and game theory. In Section 4.5, we discuss how this closure process creates legitimacy through the rhetorical strength of FTs and examine the consequences of this unique form of legitimacy.

4.2 Formalizing Technologies

In this section we discuss prior work on the use of quantification in public administration before discussing how formal mathematical frameworks have been taken up to operationalize contested policy ideals. We begin with the larger context of technologies of quantification (§4.2.1) and then define a specific subset of these, which we call *formalizing technologies* (FTs) (§4.2.2). We discuss examples of FTs and examine how they differ from other forms of quantification.

4.2.1 Technologies of Quantification

Technologies of *quantification*, which enable the "communication and production of numbers" (Espeland and Stevens, 2008, p. 401), have long been a subject of interest in the history and sociology of science. Significant attention has been devoted to the use of quantification in public administration (see Mennicken and Espeland (2019) for a more detailed review), focusing on numbers' capacity for standardization across diverse populations and for producing legitimacy.

4.2.1.1 Technologies of Commensuration

Through the process of commensuration– "the transformation of different qualities into a common metric" (Espeland and Stevens, 1998a, p. 314)– complex and disparate information is simplified into a single standardized number. In the context of policy, commensuration can reduce different interests and values to a manageable amount of information. Crucially, commensuration redefines relationships between policy alternatives by ordering them on a single scale. This ordering suggests a hierarchy between alternatives, facilitating a seemingly

optimal choice (Espeland and Sauder, 2007). Moreover, commensuration allows administrators to use the language of economic efficiency to justify policy decisions: by assigning a market price to competing interests, commensuration facilitates market-based economic solutions to value-laden policy questions (Duarte-Abadía et al., 2023). This logic has become embedded in administrative decision-making through cost-benefit analysis, which forces the commensuration of policy inputs and outputs to compare the value of policy proposals.

4.2.1.2 Technologies of Legitimacy

In addition to noting the utility of quantification for facilitating measurement and comparison, the quantification literature also examines the use of numbers to establish legitimacy. Across scientific and bureaucratic contexts, quantification has been particularly attractive in low-trust environments for its capacity to establish trust through appeals to objectivity (Porter, 1995). Trust is central in particular to legitimacy in democratic societies, where government authority must be willingly accepted by its citizens rather than imposed upon them. The authority of quantification aligns with the U.S. administrative state's historical and ongoing pursuit of legitimacy through promises of efficiency, accountability, and non-arbitrariness (Bressman, 2003).

4.2.2 Formalizing Technologies

What are formalizing technologies? Despite the rich literature on quantification, comparatively little attention has been given to variation between different forms of quantification, and different types of quantifying technologies. We conceptualize formalizing technologies (FTs) as a particular form of quantification with unique implications for public administration.

Specifically, we define formalizing technologies as mathematical frameworks that quantitatively operationalize some policy-making ideal (e.g., deservedness, fairness, efficiency, privacy preservation) and enable the design of systems that guarantee, under some explicitly stated conditions, that the outputs of the system achieve the ideal at a level pre-specified by the system's parameterization. Fundamentally, such ideals—deservedness of benefits, fairness of allocation, market efficiency, privacy—represent concepts that are value-laden and essentially contested (Mulligan et al., 2016, 2019; Viljoen et al., 2021).

Examples of formalizing technologies We can identify and group FTs by the values they quantify and the algorithmic means through which they are expressed. Consider the following examples. FTs that quantify privacy and security principles create measurements

of how much data sharing algorithms make information accessible or inaccessible to different parties as tunable functions of their parameters; example FTs include formal privacy tools like differential privacy (Dwork and Roth, 2014) and secure multiparty computation (Cramer et al., 2015). FTs that quantify equitable outcomes create measurements of inequality, unfairness, or disparity between algorithmic outputs that can be tuned with different parameters; example FTs include optimization-based modeling under fairness constraints (Mitchell et al., 2021) and algorithmic solutions to fair allocation problems (Joe-Wong et al., 2013). FTs that quantify equitable model performance create measurements of disparities in errors, variability, or stochasticity associated with algorithmic outputs; example FTs include multicalibration algorithms (Hébert-Johnson et al., 2018) and equity-aware privacy analyses (Cummings et al., 2019).

FTs represent only a small subset of technologies of quantification. FTs enact policy through formally parameterized algorithmic mechanisms for quantifying essentially contested ethical concepts. Crucially, many policy interventions have some, but not all of these properties, placing them beyond the scope of our paper. For example, quantification technologies are used to empirically measure harms within data processing systems, for instance to detect equity issues in LLM outputs or to evaluate the success of privacy attacks. Similarly, cost-benefit analysis which has been widely adopted in public policy, quantitatively assesses and predicts the net benefit of policy interventions. However, these tools lack a formal—i.e., provable, causal, ex ante—relationship between the system design and the values being quantified. Similarly, many formal mechanisms use technologies to quantify concepts whose contestation is inessential. For example, water monitoring systems may have formal properties that guarantee certain upper bounds on the concentration of lead in drinking water, but what concentration of lead is deemed safe for human consumption might be more straightforwardly negotiated through different risk-based frameworks, for instance. In contrast, whether a concept like "fairness" should be understood as equity, equality, justice, or arbitrarily many other understandings might be more fundamentally contested. Finally, FTs are necessarily parameterized, but many policy decisions are formulaic or mechanical without being parameterized; for example, allocation formulas for dispersing funding have implicit, but not explicit, parameters that dictate the policy problem's ethical dimensions.

Affordances of formalizing technologies Describing what FTs do can be surprisingly challenging. We can reasonably intuit that FTs are technological artifacts with their own politics (Winner, 1980); yet the formal logic underlying FTs produces theoretical mathematical results independent from their political context. Even asking what FTs "afford" could refer to either "how people use FTs in practice," or "how the properties of FTs admit a range of

possible uses" (Hammond, 2010; Norman, 1999). Here we refer to FT affordances in the latter sense, specifically highlighting the rhetoric of how FT developers and users engage with FTs.

FTs make explicit the algorithmic mappings between inputs and outputs. Thus these mappings can easily be made *transparent*, such that the space of decisions induced by specific FTs can be made visible to a wide range stakeholders, including participants, data subjects, and regulators. This enhanced transparency allows stakeholders to demarcate possible interventions and their inherent consequences. This transparency is apparently empowering, revealing a defined space of potential interventions.

An appealing, formalizing aspect of FTs is that they entail user-specified *parameters*—i.e. design choices—dictating how algorithms function. Because the formal reasoning underlying FTs is independent of any realized quantitative inputs or outputs to the system, FTs map the space of hypothetical inputs to hypothetical outputs. These parameters can produce constraints and locate specific policy choices—quantitative thresholds that might relate to values, such as tolerable error, trade-offs of risk or accuracy. This allows these quantitative parameters to be the locus of policy negotiation: if a value (e.g., "group fairness") were a dial, the parameters define the "notches" to which a policymaker could set the dial. Notably, FTs reflect *ex ante* choices that prescribe how algorithmic systems achieve specific quantitative properties via these notches, as opposed to *ex post* descriptions of existing systems. For example, formal privacy methods like differential privacy (DP) include pre-specified privacy loss budgets. Similarly, formal fairness methods like multi-calibration rely on multiple accuracy criteria and group definitions.

4.3 From Scientific to Ethical Closure

The process by which FTs produce legitimacy is closely linked to the process of *scientific closure*, a concept developed by sociologist of science Harry Collins (1981) to describe the social mechanisms through which scientific knowledge comes to be accepted. We argue that two forms of scientific closure, consensus and exclusivity, work together around FTs. We further argue that this scientific closure creates closure over the ethical questions that FTs purport to solve by narrowing the space of legitimate policy interventions.

4.3.1 Scientific Closure

Here we draw on two related but distinct notions of scientific closure. The first, developed in science and technology studies (STS), describes how social groups reach consensus about the interpretation of scientific and technological artifacts. In their foundational work on the

social construction of technology (SCOT), Pinch & Bijker (1984) describe the three stages of meaning-making that shape the acceptance of scientific knowledge and the adoption of technologies. During the first stage, interpretive flexibility, both the interpretation and design of a technology are malleable. In the second stage, closure, consensus develops around a single meaning and design of the technology, either through rhetorical closure or redefinition of the problem. In the third and final stage, the now-closed technology becomes embedded in the wider sociocultural context (Pinch and Bijker, 1984). This framework has been taken up and extended in the STS literature to examine the resolution of scientific controversies.

The second form of scientific closure we draw upon refers to the hidden backstage of knowledge production, in contrast to open science. From an institutional perspective, closure refers not to consensus-making but instead to limiting external intervention in science (Chubin, 1985), for example by maintaining confidentiality during peer review, delaying the public dissemination of research findings, or keeping details of the scientific process hidden. Ostensibly intended to preserve expert autonomy, this form of closure is susceptible to co-optation by other interests, for example through industry partnerships or government funding of science (Chubin, 1985). While the STS literature has dealt with issues of control over science it is worthwhile to consider this exclusivity in the context of consensus-building in order to understand how both consensus and exclusivity work together under the umbrella of scientific closure. Scientific closure, through both consensus and exclusivity, fosters legitimacy by minimizing contestation and obscuring prior disagreements to elevate certain claims as scientific. However, this legitimacy is not stable. Increasingly, waning trust in science's ability to answer policy questions had led to demands for active public participation in matters where science and technology intersect with the public interest (Nelkin, 1978).

4.3.2 Ethical Closure

In contrast to scientific closure, we use the term *ethical closure* to refer to processes that operationalize value-laden issues. Different approaches to conceptualizing and operationalizing value-laden issues may rest upon fundamentally conflicting values with social, political, and context-varying dimensions, rendering many such issues essentially contested (Gallie, 1956). To handle this problem, legal and policy interventions necessarily "constrain" ethics—specifying how values-based decisions can be engaged or intervened upon—so that institutions can act "with the force of law" (Hildebrandt, 2020). For example, legal interventions on discrimination (e.g., Title IX, re: discrimination "on the basis of sex") offer institutions such force of law (e.g., through coordinators tasked with overseeing and reporting such cases of discrimination).

We argue that, analogously, **FTs invoke ethical closure by constraining ethics with the force of technology.** While legal and policy interventions constrain how political institutions might set and enforce rules or standards, FTs leverage their technological affordances and disaffordances to constrain the space of acceptable interventions. Tools that algorithmically assess the discriminatory nature of predictions, for example, outline the space of "feasible" algorithmic configurations that trade off predictive accuracy and a quantification of their discriminatory impact (Cooper and Abrams, 2021). The adoption of a particular FT closes out policy alternatives, not on the grounds of their ethical merit, but on the grounds that they are technically incompatible with a given tool. Consequently, technical and scientific closure imposes ethical closure over otherwise open ethical questions.

4.4 Closure in Practice

The way textual policy problems are framed and technological problems are framed (Hildebrandt, 2021) necessarily has social and political effects on how they are intervened upon. However, the manner in which textual policies and FTs induce ethical closure are substantively different. We characterize how FTs create ethical closure, and the work that does to produce, shape, and even undermine legitimacy. Using examples, we show how this process that merits distinct attention from our community. We draw on two case studies in particular: formal privacy-enhancing technologies and game theoretic decision-making for military applications. We select these two cases to leverage variation in time scale, technical maturity, and the range of possible applications. Formal privacy is relatively new, and has a relatively narrow range of applications to privacy questions. Meanwhile, game theory has been deployed for decades in varying contexts, and with emphasis on a wider range of ethical problems. Despite these differences, we observe similarities across both cases.

4.4.1 Case 1: Formal Privacy-Enhancing Technologies

Privacy has been enshrined in U.S. public policy through laws like the Health Insurance Portability and Accountability Act (HIPAA), the Family Educational Rights and Privacy Act (FERPA), and the Confidential Information Protection and Statistical Efficiency Act (CIPSEA). Implementing these laws is often under the purview of federal agencies (see e.g., the eGov Act, Title 13 of the U.S. Code, and the Federal Trade Commission (FTC)'s role in privacy enforcement), with little guidance on how such implementation occurs, sometimes leading to ineffective implementations (Bamberger and Mulligan, 2010) and bouts of contestation around privacy norms (Mulligan et al., 2016). In particular, U.S. data privacy law

has historically struggled with defining the scope of "personal information" for data sharing purposes, since modern computing and data infrastructures make it easier than ever to associate data with individuals, known as the "failure of anonymization" in law (Ohm, 2009; Wu, 2013). Agencies with dual mandates to make data accessible while safeguarding privacy must resolve these issues to function successfully. Recently, agencies have looked to formal methods in privacy enhancing technologies as a promising response to this failure. We focus on policy discussions surrounding the adoption of formal privacy for the 2020 Decennial Census in this case study.

Formal privacy technologies, often referring to differential privacy (DP) (Dwork et al., 2006) and related techniques (Desfontaines and Pejó, 2019), are FTs aimed at providing ex ante measures of information leakage from data processing systems. Most formal privacy (FP) technologies use noise injection to characterize formal claims about the ability for adversaries to distinguish between results in the presence or absence of one person's data. The degree of noise injection, and thus the formal limitation on information leakage, is parameterized by what's known as a privacy loss budget (PLB); smaller PLBs refer to greater formal information leakage protections, and vice versa. In other words, PLBs associate with each data release an upper bound on the amount of information leakage that could occur. Previous work on formal privacy for the U.S. Census Bureau has examined epistemic attitudes towards Census data (boyd and Sarathy, 2022), communication practices around data privacy for data users (Nanayakkara and Hullman, 2023), and the operationalization of transparency and accountability (Abdu et al., 2024). However, little attention has been paid to the rhetoric used by the U.S. Census Bureau to justify the adoption of specifically *formal* standards.

The Census Bureau used the formal nature of FP's guarantees to justify their macroscopic decision-making process, but notably *not* their specific FP implementation decisions (i.e., what data to release and how to set the associated PLBs for each release). To explain the change from non-formal methods to formal methods, the Bureau had to empirically demonstrate that the previous methodology failed to protect data subject confidentiality. While empirical experiments correctly demonstrated that previous decennial censuses indeed posed privacy risks (Abowd et al., 2023; Abowd, 2018), no such empirical demonstrations were publicly produced by the FP system. Instead, the Bureau relied exclusively on FP's guarantees in both technical (Kifer et al., 2022) and non-technical (Keller and Abowd, 2023) communications with stakeholders. This is surprising because the Bureau implemented FP with what was considered a large PLB, implying the formal upper bound on privacy loss was too large for data subjects to infer anything useful about their particular risk; in other FP deployment scenarios, like Apple's user analytics FP deployment, such large PLBs rightfully

raise concerns about "privacy-washing" through formal, but ultimately vacuous, privacy guarantees (Tang et al., 2017). Yet regardless of these particular implementation choices, the Bureau explicitly argued that FP was the right tool for solving disclosure risk problems because of its formal properties (Jarmin et al., 2023).

This logic proved ultimately beneficial for the Census Bureau. First, it helped transparently demonstrate to the Department of Justice that the Census Bureau could ex ante satisfy its legal confidentiality mandates specified in Titles 13 and 26 of the federal code. Second, the formal guarantees ensured that no matter the background conditions by which adversarial data users might exploit Census data, the Census Bureau would not be held responsible for aiding and abetting privacy harms beyond their control. And finally, by focusing on selecting a sufficiently large PLB, the Census could demonstrate a commitment to producing useful (read: efficient) statistical estimates needed to fulfill its role managing the core statistics supporting democratic representation, appropriations for federal funding, and other official statistics uses.

At the same time, the Census Bureau's implementation left many policy questions unanswered. Many empirical conceptions of privacy harm, such as those measuring the ability to infer information about individuals (for example, race– see Kenny et al. (2021)), were labeled as "unscientific" from both the academic community (Dwork and Naor, 2010) and the Census Bureau (Jarmin et al., 2023) because they could not be simultaneously formalized (via ex ante privacy guarantees) and controlled (via choosing a PLB). To them, one could debate how to implement FP, but to suggest measuring privacy empirically outside of FP was "unscientific," "ad hoc", or not "rigorous," even when concretely tied to policy outcomes (Kenny et al., 2022). FP, in summary, helped the U.S. Census Bureau demonstrate particular commitments to both privacy protections and the production of high-quality data, but the interplay between FP and holistic policy evaluation was largely left unexplored.

4.4.2 Case 2: Game Theoretic Decision-Making and Mechanism Design

The draw of formalization is unique neither to recent technologies, nor to contemporary bureaucratic cultures. We can compare more ongoing developments in formal privacy-enhancing technologies with an earlier use of FTs to resolve value-laden policy dilemmas: the U.S.'s uptake of game theory during World War II and the Cold War. The rise (and subsequent fall) of game theory as a policy tool demonstrates both the power of FTs to enact closure, and the fragility of this closure.

Institutionally, WWII led to a major renegotiation of the relationship between science

and government in the United States. During the war, the Office of Scientific Research and Development organized significant and unprecedented cooperation between scientists and the war efforts. After the war, those same scientists were involved with developing the institutions to support science (both through federal agencies like the National Science Foundation and non-governmental organizations like RAND). The emphasis on legitimacy of these institutions was central in their founding documents, where scientific progress is framed as an essential component of government (Bush, 1945/2020), and only conducted in legitimate ways (from RAND, this must be "strictly nonpartisan, rigorously analytic" (RAND, accessed January 20, 2025)).

Importantly for our context, they elevated, and were elevated by, the adoption and legitimation of game theory in geopolitical context, which itself promised an objective, analytic way to settle political problems. In the 1920s and early 1930s, the mathematician John von Neumann and economist Oskar Morgenstern, among others, developed the mathematical machinery to capture social and psychological phenomena around games (literal games: chess, poker, baccarat). Increasing unrest in Europe after 1933 led von Neumann to imagine this formalism as a solution to the rising social instability and a tool to understand the impending international crisis (Leonard, 2010, p.189–213).

The rigor of game theory's formalizations, ironically, were a point of academic and political contestation. Concerning scientific and political legitimacy, "the insistence on 'purity' and logical rigour [was] thought necessary by Morgenstern to preserve economics from political influence" (Leonard, 2010, p.141). Yet for von Neumann, the logical rigor of this mathematics of psychology and social coalition-building was explicitly part of a larger movement of "formalist" mathematics, but also a rejection of a contemporaneous anti-formalist and pro-Nazi integration of psychology into mathematics (Leonard, 2010, p.221–223). Other scholars aligned such formal approaches with pro-democracy approaches, rather than intuition-led fascist framings, "equat[ing] the campaign for logical empiricism with the struggle for social reform" (Leonard, 2010, p.122).

As described in the STS literature, a period of interpretive flexibility precedes scientific closure (§4.3). We can view this contentious period in the American and European mathematical community as a period of *interpretive flexibility* surrounding game theory. By the 1940s, this flexibility was no more, having given way to closure. Game theoretic methods were institutionalized in U.S. policy during the Cold War, with the creation of institutions within and adjacent to the U.S. government that would bring social scientific tools to bear on policy issues. These institutions, like RAND and the National Science Foundation (NSF)'s social sciences division, brought game theory into the center of military decision-making. (Note also the central role of game theory in the popular imagination of the Cold War in

films such as *Dr. Strangelove* and *Fail-Safe*, both released in 1964.)

Erickson et al. detail what they call the "ideal type of Cold War rationality," an orientation toward formal reasoning abstracted away from context (Erickson et al., 2013). This culture of rationality reflected the U.S.'s culture of containment, creating highly controlled "symbolic enclosure[s] within which the (il)logic of nuclear politics played itself out" (Edwards, 1996, p. 118). Early rhetoric around military applications of game theory emphasized many of the same promises highlighted in contemporary discussions around formal privacy. Rather than measuring and tuning a specific value, game theory promised to integrate a host of values captured under the umbrella of "utility," or, in the Cold War context, of "military worth" (Erickson, 2019). Boundaries between science, values, and policy were consistently negotiated, with NSF authorities alternating between rhetoric that emphasized the independence of basic science from political pressure on the one hand and rhetoric that emphasized the political relevance of science on the other (Kleinman and Solovey, 1995).

Despite the enthusiasm around game theory's utility, however, critiques from inside and outside of the game theory community grew over time. These critiques mirror several critiques of contemporary FTs, including formal privacy: that they measured the wrong thing, or ultimately failed to capture the real-world impacts that people cared about. Beyond the critiques themselves, it is worth examining the varying responses to these critiques. As the discrepancy between formal and empirical became apparent, different groups took different approaches. One approach involved ignoring this complexity and treating formal models as heuristic rather than reflective of reality. An alternative approach involved making models more complicated to reflect the complexity of the empirical world. Most notably, Herbert Simon introduced the idea of bounded rationality, relaxing the more assumptions of strictly rational actors in favor of a model that tried to capture the limits of people's decision-making capacity. Building on this, Daniel Kahneman and Amos Tversky brought insights from psychology to try to capture the ways that cognitive biases might make people likely to deviate from rationality in practice. While their descriptive findings undermined the empirical accuracy of formal frameworks, Kahneman and Tversky maintained that, normatively, the rational ideal was something to strive for. By highlighting "gulf between the normative and the descriptive"(Erickson et al., 2013), their work on cognitive biases bifurcated the science of description from the ethics of normative decision-making. Yet, other factions eschewed formal modeling entirely, focusing on actionable predictions rather than generalizable theories. This emphasis on management, rather than understanding, (Rohde, 2023) reflects a move toward instrumentalism in service of efficient administration.

Game theory tried to create a closed world. While this closure ultimately proved to be fragile, it still had lasting consequences in scoping ethical questions and policy interven-

tions. Notably, game theory offered the foundations for a related but distinct FT that has since been adopted in administrative and everyday contexts. Auction theory and mechanism design, developed in the 1960s and 1970s, have seen more modern, recent wins in scientific consensus and practical adoption (and closure) (Alexandrova, 2006). The Federal Communications Commission, an American administrative agency, ran an auction for electromagnetic spectrum licenses in 1994 in consult with economists in this field; this practice has since been continually renewed and produces billions of dollars in revenue (Figliola and Gallagher, 2023, e.g.). The adoption of these auctions explicitly called upon legitimacy via efficiency (including such monetary rewards; see (Vaughan, 2003)); however, scholars dispute whether it is efficiency and academic rigor, or instead ethical closure and profiteering, that drove these decisions (Nik-Khah and Mirowski, 2019; Viljoen et al., 2021).

4.5 FTs as Rhetorical Technologies

The two use cases above reflect different time points in the research, development, translation, and adoption of FTs. But what they share in common is a set of relationships to the values embedded in U.S. administrative state governance. To investigate these relationships, we reframe the conversation around FTs by expanding our field of view. We claim FTs are not only technological artifacts, but *rhetorical* artifacts: they present information in such a way where its form, as much as its function, persuades stakeholders to prefer its adoption. This reframing does not detract from the technological promise of FTs in helping to solve difficult policy problems, but instead helps contextualize the interplay between *what* FTs contribute to policy discussions and *how* they contribute. By conceptualizing FTs as rhetorical tools, we highlight their performativity (Austin, 1975; Barad, 2003): their ability not only to reflect reality, but also to construct it. In this case, FTs meaningfully shape policy by framing political participation around a limited set of technical decisions.

4.5.1 Rhetorical Technologies

Viewing FTs as rhetorical illuminates the way that FTs are deployed to enable certain decisions, while constraining others. We focus on three rhetorical strategies shared by FTs illustrated in our analysis. First is the rhetoric of *impossibility*, which ties technological impossibility to policy infeasibility. Administrators here make claims that equate the scope of legitimate policy interventions with those legible to FTs. Next is the rhetoric of *formal quantification*, which centers formal, probabilistic, and general risks over others with more contested, context-specific interpretations. Last is the rhetoric of *control*, which emphasizes

FT implementation decisions and parameters as the locus of policy negotiation precisely *because* they allow for ex ante control of some policy outcomes. Administrators here make claims like "policy problems should be centered around FT implementation and parameter choices because they deterministically create cause-effect relationships between policy decisions and outcomes."

First, the impossibility results that motivate FT development and adoption affect the scopes in which policy interventions are deemed legitimate by administrators. In attempting to cull to broad space of possible policy interventions into something more manageable, FTs offer a seductive solution: FT researchers and developers can work on the "science" of policy mechanisms, and simultaneously, policymakers need only decide how to best deploy FTs. In our two examples, the science of policy mechanisms means the spaces of privacy-enhancing technologies with equivalent formal information leakage guarantees, or game theoretic mechanisms with the same formal equilibria properties. Administrators then have an apparent solution to the "paradox of scientific authority" (Bijker et al., 2009) by creating a clear division between the consensus-driven, scientific, and supposedly objective space where FTs are developed and the contestation-driven, subjective policy space where FTs are implemented.

As a rhetorical strategy, the division provided by FTs is highly persuasive to administrators because it adds technocratic weight to the policy alternative scoping process with seemingly minimal modification. When administrators are considering policy alternatives not immediately reconcilable with a particular FT, they're offered two paths of least resistance: either modify or translate desired policy alternatives to make them compatible with FTs, or exclude them as technologically infeasible. The ex ante nature of formal possibility and impossibility persuades administrators that, because we cannot make formal claims about particular policy outcomes, nothing of value is lost in this scoping exercise. Instead, comparisons between policy alternatives on the same formal axes should theoretically enable greater construct validity when comparing FT implementations.

Next, the logic of FTs values specific approaches to quantifying harm: formal risks, i.e. probabilistic risks over randomness inherent to mechanism implementation, are deemed more scientific based on how successfully they abstract away the background assumptions needed for making formal claims, especially when such assumptions may be difficult or impossible to verify in advance. For example, FP methods typically rely on minimal to no assumptions about a hypothetical adversary, which is seen as valuable because such assumptions involve subjective exercises of justification. Therefore for FTs, the fewer assumptions needed, the more general, universal, or scientific a particular formal claim can be. For example, formal privacy guarantees require weak or non-existent assumptions about what adversarial data users know about individual data subjects. Similarly, auction mechanisms require weak

or non-existent assumptions about the kinds of information asymmetries that market participants may have. In doing so, FTs clearly delineate between risks that are precisely, mechanistically controlled versus those that must necessarily be entangled with unverifiable assumptions or unknowable future conditions.

Administrators tasked with decision-making under uncertainty see this strategy to risk management as attractive because FTs draw our attention to risks that reduce the administrative burden needed to make high-uncertainty decisions. This reflects a general trend in how the administrative state increasingly makes decisions not about the probabilities of harmful events, but the more possibilities of these events at all (Amoore, 2013). FTs generality provides assurances that are *independent* of these possibilities, meaning the low-probability but high-risk events of interest are provably covered in this new regime. In both our examples, FTs intentionally operate without making assumptions about unknowable background information asymmetry conditions, be they about adversarial prior knowledge for privacy-enhancing technology (PET)s or market signals for game-theoretic auctions. This creates the conditions for quantitative risk management that needn't concern itself with possibilities beyond its scope or control.

Lastly, FT ex-ante guarantees are both technically and rhetorically causal, a property lacking from an overwhelming majority of technical claims made about contested ethical values. Because FTs construct their own cause-effect calculus, the mere presence of avenues for control helps to justify prioritizing the resulting metrics that get controlled, regardless of how relevant such metrics are to policy.[2] Privacy loss as a mathematical construct, for example, can be precisely bounded through careful data publishing mechanisms unlike other technologies attempts to limit disclosure risks. However, privacy loss is only one of many metrics to measure data processing's privacy harms, and the connections between privacy loss and more outcome-oriented privacy metrics is much more ambiguous. It is precisely the ability to ex-ante bound privacy loss that motivates making such a metric the locus of negotiation for policy decisions, even if it requires significant contextual interpretation.

Administrators see value in prioritizing these ex ante metrics over others because they

[2]A contemporary of von Neumann and Morgenstern offered a critique of the self-limiting rhetoric of FTs. In correspondence from Herbert Simon to Morgenstern, Simon characterizes the threat of the aesthetic, formal choices of game theory to drive their empirical relevance. In 1945 in response to their seminal work on the *Theory of Games and Economic Behavior*, Simon wrote (excerpted in Leonard, 2010, p.260):

> "I am not certain that I understand Professor von Neumann's position on this point, but I got the distinct impression from discussion with him that his preference for the definition you used was based largely on aesthetic and formal grounds. Being a social scientist rather than a mathematician, I am not quite willing for (sic) the formal theory to lead the facts around by the nose to quite this extent."

can transparently demonstrate the consequences of their interventions and their deterministic impact on policy under uncertainty. Just as with formal quantification, FTs make the managerial task of demonstrating impact trivial by drawing attention towards metrics with precise avenues for control. Contrast this with other policy interventions that require intensive program evaluations, like social benefits programs, for which their quantitative effectiveness can be more readily contested.

4.5.2 Beyond Law, Beyond Quantification

Prior work has discussed the rhetorical power of law, technology, science, and quantification. A rhetorical study of technology makes visible the reconfiguration that both technology and rhetoric create– creating new connections that reflect, enable, and constrain various forms of agency (Lynch, 2013). As technologies that formalize ethical principles, FTs act first and foremost as rhetorical devices that enact closure over open ethical questions. Both law and quantification have historically played similar rhetorical roles, supporting closure in different ways. FTs go beyond both law and quantification to combine formalizing effects of both. As rhetorical tools, FTs synthesize the strengths of quantification and law to enable closure of open policy problems and contested, value-laden decisions. This rhetorical strength is helpful for advancing bureaucratic goals, aligning with traditional pursuits of legitimacy via non-arbitrariness, efficiency, and accountability

Law enacts closure by creating certainty, with foreseeable consequences for specific events (Hildebrandt, 2020). Similarly FTs create closure by creating policy instruments where inputs map predictably onto outputs. Simultaneously, FTs borrow their force from the rhetoric of quantification. In their work on the rhetoric of statistics in political discourse, Brendan Lawson and Matthew Lovatt argue that statistical arguments function by eliding the circumstances of their production (Lawson and Lovatt, 2021). By omitting the assumptions and people involved in their design, statistics appear to be factual representations of the natural world, rather than human constructions. We observe this phenomenon in FTs, where abstraction allows administrators to elide the messy work of justifying their assumptions.

Across both the PETs and game theory cases, we observe the same pattern of closure. Returning to the closure as consensus formation (recalling the STS concept where we see closure through rhetorical closure and problem definition), we observe that in both cases, FTs 1, leverage the rhetoric of optimization and non-arbitrariness to resolve deliberation and 2, redefine the policy problem to fit within the scope of the FT's parameterization. Moreover, through closure-via-exclusion, we can see that FTs, in fact, constrain modes of participation in policy-making by limiting outside intervention. However, this closure is more fragile than

it may seem. We argue that FTs enact a form of closure based not via democratic consensus, but instead via problem redefinition and rhetorical exclusion. Yet, even when the limits of FTs are exposed, these technologies can have durable effects on policymaking approaches. As we observed in the game theory case, technologies not only have the power to change social relations but also naturalize these new social relations such that the process by which the technology was created and adopted disappear (Akrich, 1992).

Like law, FTs create closure over open ethical questions. This makes both tools appealing for policymaking, allowing decisions to be made in the face of uncertainty or ambiguity. Despite their similarities, however, law and FTs are not interchangeable. Attending to the differences between these two closure mechanisms is instructive for two reasons. First, we can identify differences in how closure is achieved and, second, we can distinguish between the policy consequences of using law or FTs to intervene in ethical dilemmas.

In liberal democracies, the authority of law and of regulation are legitimated through processes that introduce checks and balances. These checks and balances include democratic elections, judicial oversight, and public participation, for example through the notice-and-comment system. These legitimation tactics emphasize the values of democratic pluralism and create closure primarily through identifying or building consensus across a wide variety of interest groups (Rich, 1987). Moreover, via the legal system, law is legitimated through argumentation and justification, rather than through appeals to objectivity (Saunders, 1994). However, this strategy is reversed in the case of FTs. Unlike the law, the authority of FTs, is not often subject to deliberation or to checks and balances, and must therefore be legitimated by other means.

We argue that FTs, unlike law and regulation, circumvent the process of consensus building and instead enact closure primarily via rhetorical exclusion. FTs move the consensus-building component of closure to the scientific community and outside of the ambit of policy discussion. This creates closure by preserving the autonomy of technologies and the agency officials involved in its adoption to interpret the meaning of ambiguous ethical questions. For example, the Census Bureau advocated for DP by equating this particular definition of privacy with the confidentiality requirements laid forth in Title 13 of the U.S. Code. Similarly, the technical impossibility results that motivated FP development are being used to advocate for technocratically scoping the definitions of personal information in omnibus privacy laws (Altman et al., 2021). These interpretations use the perceived objectivity of scientific consensus to stand in for consensus among the broader set of stakeholders implicated in the policy process. However, the scientific paradigm for consensus-making, unlike political paradigms for consensus-making, is not built on compromise and negotiation but on expert judgment and closed institutions like peer review. Moreover, scientific consensus cannot

meaningfully stand in for consensus on the ethical issues implied by policy choices. Yet, the meaning of privacy is heavily contested (Mulligan et al., 2016). Despite widespread contestation, DP rhetorically takes consensus about the meaning of privacy as a given by appealing to consensus within a small scientific community. Consequently, FTs close off contestation about translation and operationalization decisions to the larger public and stakeholders that would traditionally be involved in the policy process. Instead, participation and contestation for this larger group is limited to a narrow set of implementation decisions.

Formalizing technologies extend the legacy of quantification in establishing legitimacy. Similarly to other forms of quantification, formalization leverages the perceived objectivity of numbers in order to limit external contestation (Porter, 1995). However, the literature on quantification has largely overlooked the different forms that quantification can take. Formalization represents a particular strand of quantification, which has consequences beyond other quantitative technologies. Green and Viljoen describe three specific features of formal thinking: emphasis on objectivity, internalism (i.e. excluding considerations that are illegible in a formal system), and an orientation toward universalism (Green and Viljoen, 2020). These features are useful for bureaucratic legitimacy: they limit the set of policy implementation choices in an apparently non-arbitrary manner. However, these features also make formal technologies particularly hard to contest: mathematical claims become unfalsifiable, both qualitative and competing quantitative forms of evidence are rendered out of scope when FTs change the status quo interpretation of policy problems, and universal negative results constrain the space of algorithmic imagination (Green, 2022).

We claim that FTs distinguish themselves from other forms of quantification through their emphasis on ex ante guarantees rather than empirical measurement. These features shape the process of closure by shaping where consensus occurs. Prior work in STS shows that scientific expertise often does not produce consensus but instead becomes an object of debate in itself (Jasanoff, 1986). However, because FTs are not tied to empirics, they are not easily contested on the grounds of evidence. When game theoretic models failed to line up with empirical evidence, it was unclear how to diagnose the problems, with many defenders of game theory arguing that solving these discrepancies was only a matter of creating more complex models.

FTs can still be susceptible to controversy. For example, the Census Bureau's adoption of differential privacy became an object of significant controversy, in part because of value-laden issues like the political role of census data and the meaning of privacy (boyd and Sarathy, 2022; Nanayakkara and Hullman, 2023). However, such contestations are rendered out of scope by the rhetorical closure that formalization enacts. Through its capacity to produce closure, FTs force debates to be on their own terms. Even when controversies emerge in the

broader public, FTs align squarely with broader administrative trends toward quantitative demonstrations of legitimacy. The adoption of FTs amidst public concern reaffirms that performances of legitimacy are often for and by the state, rather than designed to foster public trust (Barker, 1990). Specifically, the process of closure enacted by FTs closes off contestation to a small group of experts in order to preserve agency autonomy above democratic accountability. In doing so, FTs instrumentalize bureaucratic goals over other purported values, emphasizing efficiency as the key to legitimacy. While the privileging of bureaucratic goals is appealing for preserving agency autonomy, the brittleness of closure-by-exclusion can result in sudden shifts. In 2026, the Department of Commerce issued an order that would effectively ban the use of DP and related technologies for the Census and the department's other statistical product. When justifying the deployment of DP, the Census Bureau highlighted generalizability and risk-reduction. This order, by contrast, emphasized values of objectivity and accuracy, reflecting traditional administrative appeals to non-arbitrariness, and indicating the weakness of ethical closure relative to broader understandings of (and attacks on) administrative legitimacy.

4.6 Conclusion

In this paper, we show that FTs invoke ethical closure by constraining ethics with the force of technology. Specifically, FTs bypass the democratic consensus-making process by limiting the scope of contestation and those who are able to contest to more closely align with expert-driven scientific consensus-making. Relying on scientific consensus to stand in for democratic consensus forecloses the possibility of negotiation which is key to resolving ethical ambiguities in a pluralistic society. Extending beyond law and beyond mere quantification, FTs leverage the rhetorical power of both to constrain the space of legitimate policy participation and intervention. The rhetoric used to justify FTs has consequences for how policy decisions are scoped (i.e. which alternatives are legitimate or not) and how policy alternatives are evaluated (i.e. how to assess which alternatives meet particular ethical end goals). As John Lynch and William J. Kinsella write, "Rhetorical studies of technology (or technologies) ideally encourage us to identify and reflect upon the moments of decision in technological development writ large. They ask us to examine the choices that we have made during the creation and dissemination of any given technology, which we can hopefully revise or redesign." (Lynch, 2013) By reflecting on the rhetorical power of FTs, we recover the hidden political decisions embedded within the work of formalization. Through the process of formalization, policymakers crystallize new relationships between the technical and the social, between science and policy. Viewing FTs as fundamentally rhetorical matters because

the line between science and policy is inescapably blurry; we should not sustain illusions that FTs can perform this bifurcation for us (Bijker et al., 2009). In order to strive for holistic policy evaluation, we must sometimes confront evidence that resists formalization (whether formalization in general or the formalization *du jour*), even when such evidence threatens the legitimacy of FTs. Through productive contestation around ethical values, administrators might endeavor to create closure through democratic consensus-building rather than rhetorical exclusion.

CHAPTER 5

Error, Fraud, and Fodder for the Accountability Machine: The Negotiation of Administrative Error Metrics in SNAP[1]

5.1 Introduction

Through the case of the Supplemental Nutrition Assistance Program (SNAP), this paper examines how quantitative metrics govern accountability, disentangling the different components of quantification that shape how accountability is enacted. SNAP is a federal food assistance program administered at the state level. Each month, an average of over 40 million people (or 12% of the U.S. population) rely on SNAP benefits to access food. Yet, SNAP has been the target of political scrutiny, inviting increased scrutiny and anti-fraud measures. At the center of these political discussion lie a single statistic: the SNAP payment error rates (PER) – official state-level measures of over- and underpayments of benefits to households.

In a state House Oversight Committee meeting Michigan Rep. Jason Woolford argued that "The SNAP Error Rate is not some complicated idea or some myth that we're talking about today. The SNAP Error Rate is literally calculated based on MDHHS [Michigan Department of Health and Human Services]'s ability to do simple math correctly." This "simple math" supports numerous political claims and arguments about accountability, ranging from program-wide failures to inevitable human error to individual accusations of fraud. In this case, Woolford proceeded to assign blame to MDHHS director Elizabeth Hertel, suggesting she would have been fired in business world and saying that she personally cost the taxpayers this money. "Either you don't care, you're incompetent or you're facilitating fraud."

The payment error rate serves as the program's dominant accountability metric, carrying federal financial penalties that shape state administration and rhetorical weight in policy debates. This single metric governs decisions about state and individual access to resources,

[1]This chapter is based on ongoing work in collaboration with Abigail Jacobs and Simone Zhang.

accountability practices, and data collection requirements. In this paper, we examine the malleability of the payment error rate through policy changes surrounding its construction, implementation, and use, highlighting the political consequences for SNAP and the households that rely on it.

Potential fraud and waste in the administration of social services has long served as a flashpoint for political rhetoric and control of SNAP and other public benefits programs including Medicaid, housing assistance, disability support, and unemployment services. As a result of these anxieties, public benefits programs are subject to a number of accountability measures, including federal oversight, detailed reporting requirements, and extensive quality control systems. Among these measures are data collection, data sharing, and data processing: turning to numbers promises not only to identify fraud and waste in the public benefits system, but also to facilitate accountability for these problems. These metrics have largely been unquestioned, seen as neutral accounting practices and mere infrastructural underpinnings for the more political work of government oversight. Yet, in practice, accountability metrics permit a range of interpretations, support varied political claims, and shape the implementation of accountability in practice. This paper highlights how metrics can be altered to reconfigure accountability relations. We argue that error rates meaningfully shape how administrative accountability is conceptualized, and discuss how specific notions of accountability become institutionalized through data systems.

Like many other forms of quantification, these metrics have been used to legitimate and produce trust in public administration, supporting claims about how "SNAP has one of the most rigorous quality control systems in the federal government" (Food and Nutrition Service, accessed September 2, 2025). Simultaneously, however, error metrics are deployed to undermine the legitimacy of public programs and introduce further accountability and oversight. In March of 2026, Donald Trump issued an executive order "Establishing the Task Force to Eliminate Fraud," citing "unacceptable error rates," in spite of evidence that only a small proportion of errors result from fraud. SNAP error rates are permit a broad range of interpretations, in part because they conflate causally and morally distinct types of errors (see Chunn and Gavigan (2004)). SNAP error rates comprise different parties being responsible (caseworkers, families, machine error); different levels of intentionality (accidental error, intentional program violations, or systematic administrative burdens that create conditions conducive to mistakes); and different protocols and calculations across different states and different years. The direction of error (over/under) invokes distinct accountability concerns: overpayment is suggestive of fraud and waste, underpayment of barriers to access. Amidst this ambiguity, the practical meanings that influential actors attach to error rates are critical to understanding the accountability relationships that these metrics engender and

the prescriptions for reform, such as increasing surveillance and procuring software solutions, that gain traction.

Through an analysis of policy changes in the SNAP Quality Control system, this paper examines the following questions: how do politicians, advocacy groups, and administrators employ and make sense of error rates? What kinds of claims do error rates help actors advance and whose authority do they bolster? What role do data systems play in these discourses in resolving ambiguities and delimiting the space of plausible solutions?

To address these questions, we introduce a typology of action mechanisms tied to error metrics. We show that the relationship between error metrics and accountability mechanisms can be changed in three ways: by acting on operationalization, or the way error is defined and calculated, by acting on the auditing system, or the set of procedures, policies, and organizations that govern how error rates are used; and finally, by acting on incentives, or the benefits and consequences tied to error metrics. These three action mechanisms mobilize error metrics in distinct ways and, in doing so, meaningfully reshape accountability.

The remainder of this paper proceeds as follows. First, we provide some background on accountability in public administration, providing context on the proliferation of quantitative metrics to enact accountability in public benefits programs. Second, we introduce our study setting: the Quality Control (QC) System, which assesses state performance in administering SNAP benefits. We describe the error rates captured by the QC system and situate these metrics in the context of administrative accountability metrics. In section 5.4, we introduce a typology of levers that shape the relationship between metrics and accountability. We discuss the effects of changes to *operationalization*, *incentives*, and *auditing systems* on accountability relationships within the QC system. In section 5.5, we examine how these levers have shaped legitimacy in public administration and provide broader insights for the quantification literature. Finally, we conclude with a synthesis of our findings and some thoughts on the importance of error metrics as algorithmic fraud-detection systems proliferate in public benefits programs.

5.2 Background

The implementation of administrative accountability in practice reveals conflicting accountabilities: who is accountable for what, and to whom. In this paper, we examine the levers through which metrics shape accountability. This discussion contributes to the existing literature on accountability and quantification, but also provides important grounding for emerging conversations about government algorithms. There has been significant attention devoted to the challenges that algorithms pose to traditional government accountability (see

e.g. Kroll et al. (2017); Wieringa (2020); Busuioc (2021); Cooper et al. (2022b)). Yet, many of these algorithms are built on existing forms of quantification: the SNAP Quality Control system was introduced in 1977 to measure and track performance, and is now the subject of algorithmic modernization efforts to reduce error in SNAP administration. Choices about the QC system become embedded within algorithmic systems that claim to root out waste and fraud in public benefits administration. We highlight several types of quantification choices that shape accountability prior to any decisions about algorithmic system design.

We draw on Mark Bovens's definition of accountability, which comprises a relationship between an actor and a forum, in which the actor must answer to the forum and face possible consequences from the forum for its actions (Bovens, 2007). Bovens explicitly conceptualizes accountability as a relational process, as opposed to other conceptions of accountability as an organizational feature or personal moral agency. By emphasizing interactions between actors and forums—including explanation, questioning, judgment, and potential sanctions—this framework moves beyond transparency-focused accounts and provides a useful foundation for analyzing how accountability is constituted through relationships between stakeholders. This relational emphasis aligns closely with both contemporary scholarship and algorithmic accountability and the public benefits context, which must contend with accountability relations between multiple levels of government, technologists and system developers, the tax-paying public, and affected participants in public benefits programs.

While existing scholarship makes clear that public benefits recipients are subject to unequal privacy burdens, including invasive caseworker visits and substantial personal data collection (Gilman, 2012; Madden et al., 2017; Eubanks, 2018), the relational nature of accountability in the public benefits setting creates unique vulnerabilities for its participants. Demands for accountability against waste, fraud, and abuse are levied by the federal government against states, and yet, it is individuals who are coerced into giving up their privacy. Individual privacy and autonomy become fodder for the bureaucratic machine. Recent literature has characterized such vulnerabilities as a form of stategraft, where administrative processes serve as cover for exploitation of individuals by the state (Atuahene, 2023; Headworth, 2024). Other literature has shown how accountability gets reshaped through quantified and algorithmic accountability regimes (Espeland and Vannebo, 2007; Citron, 2007; Levy et al., 2021). We build upon these literatures to map the relationship between beneficiary surveillance and administrative accountability, focusing on how surveillance becomes negotiated, justified, and entrenched through metrics.

Accountability has increasingly been conceptualized in quantitative terms and specifically through performance metrics, including rankings, benchmarks, and scorecards (Espeland and Vannebo, 2007; Merry et al., 2015). In the U.S., this has involved requirements around cost-

benefit analysis in government agencies, the focus on standardized test scores in education policy, and the proliferation of performance measures in healthcare throughout the 20th and 21st centuries (Espeland and Vannebo, 2007). The creation of these metrics involves decisions, assumptions, and work that typically become obscured by the number itself. The sociology of quantification literature has largely focused on these myriad decisions involved in the production and interpretation of numbers and on those who make these decisions (Espeland and Vannebo, 2007; Lewis, 2015; Bosk, 2020). This literature also emphasizes the way that these metrics become institutionalized, arguing that changes introduced through quantification become deeply ingrained and difficult to disrupt (Chun and Sauder, 2022).

Quantification is particularly central to governance, both for facilitating public administration and for holding governments to account (Mennicken and Espeland, 2019). In the United States, cost-benefit analysis serves both functions– simultaneously guiding regulatory decision-making and mediating agency accountability to the courts. Beyond cost-benefit analysis, metrics serve a widespread accountability function in agencies. As Lawrence L. Martin (1998) writes in his account of the proliferation of performance metrics in U.S. government programs in the 1990s, "[a]s originally conceptualized, performance measurement was to be a structured method by which government programs would communicate with their stakeholders (e.g., elected officials, citizens, advocacy groups, clients and others)." Quantitative accountability has been particularly essential in establishing and maintaining the legitimacy of the administrative state. Metrics inform how agency officials make decisions about how to administer policies with multiple or ambiguous goals. Agencies may use them to track whether programs deliver services efficiently, apply eligibility rules accurately and consistently, prevent improper payments, and reach intended beneficiaries. They simplify complex information, facilitating insight into the intricate web of institutions and agency officials that make up the administrative state (Espeland and Vannebo, 2007). Finally, the mandatory public circulation of these metrics promises to meet agencies' democratic imperatives by enabling public scrutiny (Martin, 1998).

Metrics are especially salient in decentralized administrative systems, where oversight authorities like the federal government seek to evaluate and steer the activities of dispersed implementing units. Many U.S. public benefits programs are administered through layered arrangements, with responsibilities distributed across federal agencies, state agencies, and local-level administrators, contractors, and frontline workers. In such complex governance structures, performance metrics can be attractive because they offer a simple, seemingly politically neutral means of governing from afar, even as these settings are precisely where metrics may have the most distortionary impacts (Moynihan et al., 2011). Amid complexity and variation, quantification projects like error rate calculation systems offer a means to

translate heterogeneous case-level decisions into aggregate measures of performance. In so doing, they commensurate implementing units, making it possible to define thresholds of acceptable performance and giving oversight authorities a concrete object to which consequences can be attached.

By facilitating comparison, error metrics can be politically consequential. They can obscure choices – what counts as an error, how cases are sampled and reviewed, how evidence is collected and weighed – that render distinct decisions, actions, and contextual conditions functionally equivalent (Espeland and Stevens, 1998b). As error rates aggregate caseworker mistakes, procedural lapses, household reporting errors, documentation gaps, and the routine mistakes that administrative complexity breeds, the resulting number does not specify what kind of failure has occurred. Distinct problems are collapsed into a single error rate that is then subject to interpretation, including by political actors who can employ error as evidence of waste, fraud, or abuse. Such interpretations may be particularly concerning as fraud-centered narratives can place recipients as the most visible subjects of accountability. Much as public access to criminal legal data has been justified as a means of institutional oversight despite primarily serving to enable the surveillance and stigmatization of individuals (Lageson, 2020), error-based accountability systems can heighten burdens for beneficiaries, who must shoulder the surveillance, documentation demands, verification requirements, and risk of lost benefits they bring (Press and Tanur, 1991; Headworth, 2021).

In this paper we build on the sociology of quantification literature to understand how accountability metrics can be intervened upon. While official error metrics may seem settled, they depend on and are sensitive to a number of choices. In this paper, we excavate some of these choices, highlighting important changes in the QC system and tracing how these changes reshape accountability relationships. We show that metrics systems encompass distinct structural components that can be acted upon with distinct implications. These components change what policy outcomes are seen as commensurate, what political claims can be supported, what solutions are considered viable, and what accountability relations are emphasized or neglected. In the sections that follow, we present a typology of these metrics components and illustrate why they are useful to examine both individually and in concert.

5.3 Accountability Metrics in Practice: SNAP Quality Control

We examine the relationship between metrics and accountability in the context of the SNAP Quality Control (QC) system. SNAP, formerly known as the Food Stamps Program, is the largest U.S. food assistance program, serving 41 million people per month. SNAP provides means-tested benefits to eligible low-income recipients, who can use these benefits to purchase eligible food items at over 260,000 authorized retailers (CBPP, 2024). Concerns about fraud in SNAP have featured prominently in discussions of the program since its inception. In recent decades, the deployment of algorithmic, data-driven systems to automate fraud detection has grown increasingly important to efforts to enhance accountability in the program (GAO, 2018). Even prior to the adoption of contemporary algorithmic fraud-detection systems, the use of data to govern the integrity of SNAP administration has played out through the SNAP Quality Control (QC) system.

Program integrity and fraud detection in SNAP administration are documented through the quality control or QC system. The QC system is designed to assess the accuracy of benefits payments and eligibility determinations within the program. State agencies responsible for administering SNAP conduct their own reviews. A sample of these reviews are subsequently verified at the federal level by United States Department of Agriculture (USDA), who uses this data to determine state and national error rates. In the event that a state's payment error rate is determined to be excessive, the state must either pay back the excess directly to USDA or invest half of this amount in addressing the source of the errors (Food and Nutrition Service, accessed September 2, 2025).

Each year, a sample of the data underlying the QC system are made public. This data is produced annually based on a sample of households selected for review as part of the SNAP QC system. Both the system and the data cover reviews conducted by state SNAP agencies across all 50 states, Guam, the U.S. Virgin Islands, and Washington D.C. The QC data govern the accountability relationship between states and the USDA, and are also used by government internal government oversight bodies like GAO to determine whether USDA is managing the program appropriately (U.S. Government Accountability Office, 2024a) and by external policy organizations, like the Urban Institute and Brooking, for policy analysis (Bauer and Whitmore Schanzenbach, 2025; Gutierrez, 2025). Thus, this data underlies accountability relationships between state and federal government, between different branches of the federal government, and between the government and civil society.

While we focus on the SNAP QC system in this paper, similar systems recur across many decentralized benefits programs. Some of the most prominent of these were built into federal

means-tested benefits programs. Following the Food Stamp Act of 1977's formalization of the SNAP QC system, analogous QC systems were also in operation by the early 1980s in the Aid to Families with Dependent Children (AFDC) and Medicaid programs. All shared a common architecture, consisting of a sample-based review of case records that produces aggregate measures of error or improper payments, which were then reported, compared, and linked to penalties and corrective actions. This overall approach was later significantly expanded through the Improper Payments Information Act of 2002 and the Payment Integrity Information Act of 2019, which required federal agencies to measure and publicly report improper payment metrics in programs "susceptible to significant improper payments." As a result, error-rate systems now organize oversight and accountability across an increasingly wide array of public benefit programs. Unpacking the operations of the SNAP QC system – one of the longest-running and most institutionalized such systems – can thus offer insight into how this increasingly prevalent approach governs benefits administration and distributes the burdens of accountability across agency leaders, frontline workers, and beneficiaries.

We focus specifically on *error metrics* within the QC system. The QC system tracks multiple types of data, including demographic and economic information at the household-level. However, these data are ultimately collected in service of assessing payment errors, state liabilities (i.e. the amount that states must repay to USDA), and improve accuracy in benefits and eligibility determinations (Food and Nutrition Service, accessed September 2, 2025). Error metrics are central to these goals. Within the QC system, there are two primary error metrics: Payment Error Rate (PER) and Case and Procedural Error Rate (CAPER). PER estimates the percentage of incorrectly issued SNAP benefits (i.e. overpayments and underpayments). Meanwhile, CAPER evaluates so-called "negative cases"—cases in which benefits were denied, terminated, or suspended— to assess whether the decision was correct, whether procedures were followed, and whether households were notified in a clear, accurate, and timely manner. Unlike PER, CAPER does not measure the magnitude of improperly withheld benefits, but instead simply captures the percent of cases reviewed in which an error occurred.

We conceptualize the QC system as an accountability *infrastructure* (see (Star and Ruhleder, 1996)), which supports accountability relationships in the administration of SNAP while remaining largely unscrutinized. Specifically, we draw on Mark Bovens's definition of accountability in public administration as a social relation (see §5.2. Bovens's account emphasizes three central features: 1) a *relationship* between an actor and a forum, in which 2) the actor must be able to *explain and justify* their conduct to some forum, who 3) can *pass judgment and enact consequences* for the actor (Bovens, 2007). In the context of algorithms,

the relationship between actor and forum is complicated by the network of actors involved in the design, implementation, and use of algorithmic systems (Wieringa, 2020). A relational approach to algorithmic accountability emphasizes analysis of the power dynamics between various actors within this complex network, explicit scaffolding for public input (Metcalf et al., 2021), and prioritizing those who are disproportionately impacted by sociotechnical systems (Birhane, 2021). A careful analysis of who is assumed to be accountable, for what, to whom, and under what circumstances (Cooper et al., 2022b) can help shed light on the accountability relations supported and/or undermined by the SNAP QC data infrastructure.

In the following section, we outline how changes to the QC system shape accountability, often while evading scrutiny. The United States Department of Agriculture Food and Nutrition Service (FNS) formalizes changes to the QC system through updates to FNS Handbook 310, which outlines procedures for conducting QC review. This update and other operationalization changes are issued as memoranda, which do not undergo agency rule-making procedures.[2] These rule-making procedures are a central tenet of accountability in broader administrative decision-making, but are not a part of the QC accountability infrastructure. Moving decisions about the QC system outside of notice-and-comment rule-making limits the accountability role of the public (i.e. makes SNAP administration less accountable to the public) and frames decisions about metrics as neutral technocratic decisions that should be left to agency experts within USDA rather than value-laden policy decisions that require broader stakeholder input. We highlight where these policy decisions occur to support greater engagement in these choices.

5.4 From Metrics to Accountability: a Typology of Levers

The literature on quantification and accountability suggests that numbers have become central to the performance of accountability in government and beyond (Espeland and Vannebo, 2007). However, quantification captures a number of moving parts: what to measure, how it is measured, what this measurement means, and how these decisions are made. Abstracting away the wide range of decisions that undergird the quantification process obscures the dynamics that shape accountability. In other words, by focusing on the number itself, critiques of quantification risk losing sight of the many places where interventions can occur. In this section, we identify and disentangle the different parts of the quantification process through

[2]In January of 2025, USDA issued a proposed rule to formally incorporate FNS Handbook 310 into SNAP regulations. If implemented, this would ensure that changes to the handbook undergo administrative procedures required for agency rule-making, including a notice-and-comment period.

the example of the SNAP Quality Control system. Through this analysis we specify the precise mechanisms that allow metrics to govern accountability.

The QC system governs accountability relationships across the entire SNAP administration process. Through this system, error metrics provide the basis for broader accountability processes. Data collection, calculation, audits, and policy around QC error metrics shape who is accountable to whom, and for what. Although the QC system has been central to SNAP for years, this system is not a fixed assessment of performance: the way that administrative error is constructed, understood, and deployed has changed regularly within the QC system's history. The error metrics and the system used to calculate them obscure a variety of policy levers that can be mobilized to reconfigure accountability relationships within the broader SNAP ecosystem. In this section, we identify three types of levers that federal and state governments deploy to change the relationship between metrics and accountability. Taking the SNAP QC system as an illustrative case of accountability metrics, we identify three types of changes that the QC system has undergone with illustrative examples for each. We do not claim that these types of levers are exhaustive or descriptive of all accountability metrics, but we show that attention to these levers can provide a framework for examining exactly how accountability is reshaped through quantification.

First, we discuss how error is *operationalized*: which accountability constructs are measured, and how they are transformed into metrics. Second, we examine how error is tied to *incentives* through policy that governs the penalties and benefits associated with metrics. Third, we consider the broader *auditing system*– the practices and procedures that surround the construction of error on the ground. Finally, we describe the *interactions* between levers, highlighting entanglements within the accountability infrastructure. As we describe, any of these levers can be altered to reconfigure who is accountable, for what, and with what consequences. For each type of lever, we compare specific policy changes in the SNAP QC system to illuminate how changes to these dimensions of error rates redefine both the values underpinning accountability and the enactment of accountability. We show that despite their stability in rhetoric about program performance, error metrics are malleable. Through changes in scope, policy, and practice, SNAP error rates are capable of both supporting multiple notions of accountability and subverting accountability altogether.

5.4.1 Operationalization

One way to change the relationship between metrics and accountability is to change the construction of the metrics themselves. In other words, which accountability constructs are measured, and how they are transformed into metrics—their operationalization—changes

what accountability dictates. To illustrate the effect of operationalization changes on accountability in the SNAP system, we examine the case of a 2022 update that changed how overpayments were calculated in the QC system.

Example 1–The operationalization of procedural errors: In 2022, USDA begins counting the entire benefit amount for cases with certain procedural errors as overpayments. Previously, the discovery of a procedural error like a missing signature would lead QC reviewers to verify whether the household was eligible and received the correct benefit amount based on its income and other circumstances. Under the new policy, this second review is not conducted. Instead, the procedural error leads to a determination that the entire benefit amount was an overpayment, even if the benefit calculation was substantively correct.

In the event of an overpayment caused by procedural error, the household must repay the entire benefit amount even if the error was the result of state error. Moreover, the policy change was not mirrored for underpayments: if households were erroneously deemed ineligible or procedure for denial is not followed correctly, these errors are not counted as underpayments. These errors make up a separate error rate called case and procedural error rate (CAPER) and are not reflected in the QC. Rather than measuring over- and under-payments, CAPER assesses only cases where benefits were denied, suspended, or terminated. For these cases CAPER captures whether procedures were correctly followed, whether decisions to deny or terminate eligibility were accurate, and whether households were notified in a clear and timely manner. Unlike Payment Error Rates (PER), which are central to the QC system of incentives, CAPER errors are not tied to federally imposed financial penalties for state agencies.

This asymmetry reflects how accountability is conceptualized and where it is considered relevant. Under the new policy, states are accountable for procedural errors that lead to benefits being allocated but not for procedural errors that lead to benefits being denied.

Example 2– The error tolerance threshold: However, changes to operationalization need not emphasize errors as the site of accountability. In contrast to the previous example, which expanded what was considered a payment error, the federally mandated error tolerance threshold narrows what is considered a payment error. Small errors–defined as payment errors lower than the tolerance threshold– are excluded from the error rate. In other words, states and administrators are not seen as accountable for minor errors, which are conceptualized as the inevitable cost of doing business. To justify raising the threshold, FNS argued that:

> To improve business efficiency, agencies must prioritize those areas that have the most potential to improve payment accuracy and reduce improper payments. This rulemaking supports that goal by focusing on errors that are the most

> economically efficient to correct. The provisions of this rule will improve the data available at the Federal level allowing for further analysis of the root causes of payment errors. The Department's payment accuracy team will be better able to focus on the largest and most problematic errors and then work with States on additional cost efficient ways to improve the Administration's goals to reduce improper payments.

A sufficiently high threshold allows states to prioritize efficient case processing rather than spending more time to verify that benefits determinations are exactly correct. By contrast, in 2025, Iowa lawmakers Senator Joni Ernst and Congressman Randy Feenstra introduced a proposal to effectively eliminate the threshold by reducing it to $0. This competing view of accountability views all errors as being within scope for allocating blame and enacting consequences, even if addressing these errors at the expense of slower processing times or reduced access.

5.4.2 Incentives

The second lever involves the network of incentives tied to metrics. Policy changes that shape the penalties and benefits associated with metrics transform accountability relationships by shaping the ability of different parties to impose consequences on others. As Bovens's (2007) conception of accountability highlights, the ability to enact consequences is essential to the accountability relationship between actor and forum. These incentives dictate who is accountable to whom: who is seen as the actor, who plays the role of the forum, who can impose consequences, and who is subject to these consequences.

In the case of SNAP, incentives play a major role in enacting accountability for benefits administration. The most significant incentives for SNAP administration are tied QC metrics, and in particular to the PER, which measures both under and overpayments of SNAP benefits. Notably, these incentives are typically tied to penalties rather than benefits. These penalties include repayments to USDA, increased oversight, and, in the most extreme cases, suspension of federal funds. As described in the House Report on the Food Stamp Act of 1977, "This penalty is intended to provide a significant financial incentive for States to improve overall program management." (U.S. Congress, House Committee on Agriculture, 1977, p. 361) In this subsection we argue that changes to these incentives alter the locus of accountability, both in terms of where blame for errors is assigned and to whom.

Example 1– Changes to state-level sanctions and bonuses: To illustrate this, we can examine the 2002 Farm Bill's changes to the structure incentives tied to the payment error rate. Prior to 2002, sanctions were assessed to all states with PER above the average. This regime

assumed that Beginning with the 2003 QC system, the focus changed to assess sanctions to states with a statistically significant probability of above average PER for two consecutive years. This change reflects a shift in how accountability is conceptualized: under the pre-2002 model, any below average performance in a given year is seen as something for which states are accountable. In practice, this meant that about half of states were penalized each year. After the policy change, however, single year performance was no longer something for which states were accountable. By emphasizing both statistical significance and a longer time period, the new policy frames short-term error rate differences as potentially random, outside of the control of administrators, or even caused by changes that support long-term improvements to public administration but require a transitional period such as new computer systems. Simultaneously, the 2002 Farm Bill offered bonuses to states with low and improved PER AND states with "high or improved rates of serving eligible households and in doing so in a timely manner." (CBPP, 2003, p. 7). These bonuses are awarded annually to the states with the top and most improved performance for four different measures: the payment error rate (PER), the case and procedural error rate (CAPER) [3], the program access index, and application processing timeliness. In contrast with the sanctions which are based only on the payment error rate, these bonuses are tied to a variety of metrics, broadening the scope of outcomes for which states are accountable.

Example 2– Introduction of state cost-sharing: Both sanctions and bonuses in the SNAP QC system are paid by states to USDA. As a result, this is the primary accountability relationship between in the SNAP system. Other accountability relationships could target USDA, individual administrators, technologies used in processing applications. They could focus on accountability to oversight agencies, non-governmental watchdogs or advocacy organizations, SNAP recipients, or the general public. However, the incentive structure and the error rates themselves both emphasize the relationship between states and USDA. Errors rates are primarily published and tracked at the state level, even in states where county governments are responsible for administering SNAP benefits. Moreover, USDA does not directly encounter incentives to keep overall SNAP error rates low.

To further illustrate how this dynamic is governed through policy decisions about incentives, we discuss the 2025 change to cost-sharing rules between state and federal government. This policy leverages newly mandated state contributions to SNAP benefits to penalize states with high error rates. Specifically, the One Big Beautiful Bill Act requires that states with error rates above 6% participate in cost-sharing for the first time, marking a shift in who

[3]In 2002, the case and procedural error rate was called the negative error rate. Both metrics measure whether terminations, denials, and suspensions were correctly determined. Beginning in Fiscal Year 2012, the name was changed to reflect the additional measurement of whether procedures were correctly followed, including for example timely notification.

bears the burden of SNAP benefits. This policy directly ties the payment error rate to financial burdens for states: the cost-sharing rate rises with error rate in a tiered system, where states with the highest error rates are required to pay 15% of benefit costs.

This change to cost-sharing further shifts the burden of accountability toward state administrators rather than the federal government. This policy change was justified in a memo encouraging increased accountability for states arguing that under the previous policy, "the SNAP benefit is 100 percent funded by the Federal government, resulting in minimal incentive for States to control costs, enhance efficiencies, and improve outcomes for recipients" (Thompson, 2025). Under the new cost-sharing rules, errors are conceptualized as originating with improper administration rather than federal policy. State-level payment error rates make states illegitimate administrators of federal funds, but high error rates throughout the SNAP system in fact decrease the federal government's responsibility for SNAP costs. Error rates often rise in response to federal policy changes, which requires administrators to adapt to new rules. Under the new cost-sharing rules, however, states are the only party accountable for these changes.

In addition to shifting *who* is accountable further toward states, the new incentives also shift *what* they are accountable for. Before the policy change, overpayments were the main source of penalties for states, who are required to repay any excess benefits payments. The disincentives around overpayments coupled with a lack of incentives around other goals like hunger, program participation, and quality of service has led administrators to pursue low error rates at the expense of access for decades (Bolen, 2001). Because the new cost-sharing incentives are based on payment error rate, which treats overpayments and underpayments equally, the policy introduces greater accountability for underpayments. By the same token, because the new cost-sharing rules are based on payment errors and *not* on CAPER, responsibility for the wrongful denial of benefits is de-emphasized. Prior to the new cost-sharing rules, both PER (which measures over- and under-payments) and CAPER (which measures wrongful denials or terminations) resulted in similar forms of accountability: a high error rate would trigger increased oversight through a corrective action. The asymmetry of the new policy shifts what states are accountable for, encouraging accountability for payment errors over accountability for under-provision of services.

We observe the consequences of these incentives through subsequent changes in state-level administration. In response to the new cost-sharing policy, state governments proposed and initiated a number of changes. These changes included a proposal to pass on incentives to individual administrators, as in OK SB 1310, which would introduce pay cuts for administrators when Oklahoma's payment error rate is above 6% and bonuses when it falls below 5.75%. Micheal Bergstrom, the state senator who introduced this bill, said: "I want

some accountability. It's never okay to waste taxpayer dollars at any level of government. So, if the error rate is above six percent the DHS director and every SNAP administrator and supervisor should take a hefty pay cut, and the higher the error rate the bigger the cut." (Oklahoma Senate, 2025) This vision of accountability ties broad state-level metrics to individual administrators, triggering automatic penalties regardless of the source of error. Notably, this bill mirrors the 6% threshold set at the federal, demonstrating how one seemingly arbitrary policy change to incentives can have widespread consequences on how accountability is implemented in practice.

Other state-level changes in response to new incentives echo elsewhere in the QC system, with multiple states, including Georgia, Mississippi, and New Hampshire, calling for increased data collection or data sharing. These measures range from citizenship verification to more frequent recertification to data sharing agreements across state agencies to verification with external data sources like credit bureaus (Pattinson-Gordon, 2025; Kallins and Thompson, 2026). The resulting privacy reductions for SNAP participants are rendered insignificant by the focus on error rates as the primary target of accountability. Similarly, reductions in program access are seen as acceptable ways to lower high error rates. The most stark of these changes is apparent in Arizona, where SNAP enrollments declined rapidly due to new, onerous requirements— the state experienced a 47 percent drop in participation in the first seven months following the new cost-sharing rules despite a state-wide increase in unemployment (Santa Cruz, 2026). State officials argue that this drop in participation is due in part to efforts to keep the payment error rate below 6% while experiencing staffing shortages and increased administrative costs. Under the new cost-sharing structure, both states and households are punished for high payment error rates.

Example 3– Eliminating USDA's discretionary powers: Through both the 2002 and 2025 changes to incentives tied to error metrics, we note that changes to the incentive structure have resulted in increased accountability for states, individual state administrators, and public benefits recipients, while federal policymakers have largely played the role of holding accountable (as opposed to being held accountable). Yet, even the accountability role of federal government is malleable under changes to the SNAP incentive structure. The 2014 Farm Bill eliminated USDA's discretionary powers to waive any portion of a state's QC liability. Prior to 2014, the secretary of agriculture could waive SNAP penalties assessed to states at their discretion. This policy provided flexibility in cases where states encountered higher than normal error rates due to short-term circumstances like new policies, new technologies, or economic shocks. The elimination of this discretion emphasizes uniform application of rules, as opposed to the context behind numbers. This conception of accountability emphasizes consequences for outcomes over process. This shift also seeks to

eliminate the forum to whom states are accountable. Recall Bovens's definition of accountability as answerability (including a duty of explanation) to a forum who can pass judgment and impose consequences. Traditionally in a public administration context, the tax-paying public might inhabit the role of the forum, or federal oversight might occupy this role on the public's behalf. Not only is the public largely absent from accountability relationships under the QC system and incentive structure, but the elimination of discretion minimizes even the federal government's ability to pass judgment. While it may seem that states are accountable to USDA, the turn to quantification renders USDA's judgment irrelevant in favor of metric-driven rules.

5.4.3 Auditing System

The third type of change to the accountability system involves changes in practices surrounding error rates. While operationalization dictates how error is defined, the broader *auditing system* determines how agency officials approach the task of measuring error. This auditing system comprises both the formal *policies* and informal *practices* that guide state agencies' approach to measuring program performance and integrity. This includes auditor training, quality control sampling decisions, documentation requirements, and procurement policy, as well as the everyday practices that auditors use on the ground. Many of these auditing practices are informal and documented primarily through ethnographic work. Spencer Headworth (2021) for example, details the folk theories that underpin fraud workers' decisions about what constitutes SNAP fraud. Beyond these informal practices, other changes to auditing are systematized and widely adopted through policy or top-down guidance.

We highlight two key decisions the SNAP QC auditing system: one dictated by both federal- and state-level policy, the other by USDA guidance. In both cases, changes to the auditing system not only affect the reported error metrics but also reveal what forms of accountability are seen as most meaningful in practice.

Example 1– Interview requirements: The first policy decision involves the interview mandated through QC review. Federal guidance requires a face-to-face interview for all households receiving over $100 in benefits. The original spirit of this guidance highlights a version of accountability based in suspicion and physical verification of recipient circumstances. Notably, the FNS Handbook encourages vigilance and observation of the person's home, like evidence of unexpected children's toys (U.S. Department of Agriculture Food and Nutrition Service, 2025b). However, the requirements have become looser, particularly since COVID-era waivers that allowed states to conduct reviews by video or phone in certain circumstances. This change, contrary to some of the other examples we have discussed, enables a shift to-

ward more efficient administration and greater accessibility. Even under the new interview rules, federal guidance maintains an emphasis on observation, as in the following hypothetical example: "A household consisting of a mother and three young children is completing a video/telephone interview for QC purposes. During the interview the reviewer hears a man's voice in the background. This may be an indication of another household member in the home. The situation would warrant further inquiry" (U.S. Department of Agriculture Food and Nutrition Service, 2025b). In practice, enabling increased flexibilty in state-level policies about how to administer interviews allows states to determine their own balance between suspicion, efficiency, and access.

Example 2– Fraud Framework and Implementation Grants: USDA released the SNAP Fraud Framework in 2018 to provide states with a flexible strategy for fraud prevention, detection, investigation, and overall performance management. The framework is organized around seven components: organizational management, performance measurement, recipient integrity education, fraud detection, investigations and dispositions, analytics and data management, and workforce development. To encourage implementation, FNS has funded state-level SNAP Fraud Framework Implementation Grants since 2019, supporting state investments in data analytics systems, fraud detection technologies, and staff training. Recent grant solicitations have explicitly funded projects involving predictive analytics and automation (U.S. Department of Agriculture Food and Nutrition Service, 2025a). The Fraud Framework and associated grants have important implications for state-level auditing systems. Rather than relying primarily on traditional case reviews or annual QC audits, states are increasingly encouraged to adopt continuous monitoring systems that identify potential risks throughout a recipient's participation in SNAP. As a result, accountability is no longer conceptualized as periodic verification of eligibility but as ongoing surveillance and risk management, as we will further discuss in §5.5.

5.4.4 Interactions Between Levers

Disentangling the various levers that govern accountability can help illuminate important choices within quantification systems. However, the effects of these levers is rarely isolated. We provide two examples to demonstrate the utility of examining the accountability levers both separately and in concert.

Example 1– Asymmetric penalization of procedural errors: Operationalization decisions are largely not accountable to the public. Yet, they interact with decisions about incentives with important consequences. We previously described the decision to count cases with procedural errors as overpayments in the PER calculation whenever benefits are granted, but

relegate procedural errors in the case of denials and terminations to CAPER (thus excluding these "underpayments" from the PER calculation). We also noted that under the 2025 change to cost-sharing incentives, states are penalized for the payment error rate (PER) but not for the case and procedural error rate (CAPER). The combination of changes to asymmetric operationalization of procedural errors and different incentives tied to PER and CAPER results in substantially greater accountability for procedural errors tied to benefits determinations than procedural errors tied to terminations or denials. This interaction frames accountability around spending rather than service provision.

Example 2– Shelter cost verification: Since 2002, SNAP recipients have been eligible for an income deduction for shelter costs paid above 50% of their income. Because of the unpredictable nature of shelter costs and difficulty of obtaining documentation, states can opt to rely on self-reported costs rather than requiring documentation. However, if these cases are selected for QC review, caseworkers must verify the exact shelter costs, which can lead to a QC error that contributes to the PER. This mismatch between eligibility procedures and auditing procedures highlights a level of flexibility that states use to prioritize other administrative goals above error metrics. Some states—including California, Missouri, Washington, and Wisconsin—elect not to optimize their error rate in favor of more efficient allocation of benefits, especially for populations with unpredictable housing costs like recipients facing housing insecurity. However, the introduction of the new cost-sharing rules undermines this flexibility, especially for states that face costs that could jeopardize the future of their SNAP programs altogether. Since the introduction of the cost-sharing rules, multiple states—including Illinois, Massachusetts, Nevada, Pennsylvania, Texas, and Virginia—revised their eligibility policies to match their auditing policies, requiring verification and documentation of shelter costs and eliminating applicants' ability to self-attestation. These changes highlight the potential of both the auditing system and strong incentives in combination to preclude states' flexibility to set their own administration priorities and limit possibilities of experimentation or innovation in administration strategies.

5.5 Discussion

Our analysis shows that addressing error rates have become central to states' performance of legitimacy in SNAP administration. A well-documented critique of quantification is its tendency to exclude and overlook considerations that are difficult to measure (Scott, 1998; Espeland and Stevens, 1998a; Mennicken and Espeland, 2019). The case of SNAP error metrics, however, shows that even among easily quantified policy goals, specific operationalization, incentives, and auditing choices all determine the effects quantification. Although

both procedural and payment errors are measured within the QC system, policy choices have led the payment error rate to dominate as the primary accountability metric governing SNAP administration.

Despite the numerous choices involved in the construction of administrative error—across operationalization, incentives, and auditing systems— the rhetoric around error rates largely treats these figures as objective measures of performance. Despite the USDA's claims that error rates are distinct from fraud, discourses around error rates are deeply entwined with fraud in practice. Through this fraud logic, which identifies fraud as the primary source of illegitimate program administration, highly visible error metrics stand in for transparency and accountability into the broader SNAP system. By obscuring distinctions between fraud and other aspects of system management, error rates center beneficiary data collection in SNAP's accountability infrastructure. The QC error rates are themselves constructed based on two-tiered review system that involves both state-level and federal-level review of detailed household information. We argue that ambiguity in the meaning of error functions rhetorically as a moral authorization for surveillance and reverses traditional accountability relationships. Typically, accountability focused on services by the state must justify depriving recipients of benefits. However, the assumption that any error might be fraud inverts this accountability relationship and presumes that the recipient must continually justify eligibility. We propose that through this subversion of accountability, the administrative state's legitimacy has become grounded in suspicion and surveillance rather than in its ability to serve the public interest. This subversion of accountability falls unevenly on participants in social benefits programs: "Documentation of household circumstances and finances grounds applicants' official categorization as program eligible. And when fraud units challenge that categorization's legitimacy, documentary evidence becomes key in clients' incrimination or exoneration. Thus people whose lives tend to be more unstable, more stressful, and worse suited for easy documentation (and the accumulation of documentation, given precarious housing) face elevated requirements to provide it" (Headworth, 2021, p.220).

Even when the source of program errors is understood as accidental or even inevitable (Kallins and Thompson, 2026), reduced privacy for recipients is seen as necessary fodder for the accountability machine, which is built to prioritize accuracy over all other aspects of good government. This privileging of accuracy lends itself to primarily technical solutions. Accountability becomes abstract with fault mostly assigned to the inevitable error that comes with the old system and accountability is only rendered visible through data collection: "Every piece of paper that one of our eligibility workers have to touch and do some manual data entry, that's a potential for an error," says Mark Jones, chief communications officer for the Mississippi Department of Human Services, as quoted by Governing (Pattinson-Gordon,

2025). "We're trying to get as much data, we're trying to do as much training, and we're trying to do as much technical upgrades as we can."

Error rates and the incentives tied to them, are used to justify the adoption of algorithmic technologies, which promise to increase payment accuracy. The long-term sidelining of the forum in SNAP's accountability infrastructure opens the door to a system where accountability is constructed only in terms of its responsiveness to metrics, rather than to any legitimate stakeholders. Even before fraud-detection algorithms entered the picture, the SNAP QC system was configured around payment accuracy and suspicion, rather than public service. Algorithmic systems not only inherent these priorities, but amplify them by optimizing payment accuracy even at the expense of intensified data collection or reduced access to public benefits. However, this need not be the case. As our analysis shows, interventions across operationalization, incentives, and auditing open the door to alternatives. The payment error rate has changed to reflect that inevitable minor errors do not reflect fatal programmatic deficiencies or that short-term error rate increases might allow for useful innovations. Similarly, the existence of the CAPER points to error metrics that prioritize the consequences of improper denials of benefits for recipients over the cost to the government of incorrect payments. Exploring the space of possibilities, and the levers that allow us to intervene in this space, highlights the potential to reintroduce SNAP participants and the public into the program's accountability infrastructure.

5.6 Conclusion

This paper has examined how accountability is constituted through quantification by tracing changes to the SNAP Quality Control system. We show that error rates are the product of a series of policy choices concerning how error is defined, measured, attached to consequences. We therefore highlight three levers through which quantification reshapes accountability: operationalization, incentives, and auditing systems. Changes to any one of these dimensions can alter who is accountable, for what, to whom, and with what consequences. Metrics, through their construction and adoption, actively produce and redistribute accountability relationships throughout the administrative state.

Our analysis contributes to the literature on quantification by shifting attention away from metrics themselves and toward the infrastructures that make metrics politically consequential. Existing scholarship has demonstrated that quantification obscures the choices involved in translating complex social and political phenomena into measurable indicators (Espeland and Stevens, 2008; Deringer, 2018). Other work attends to the way that these numbers become institutionalized, often remaining consequential even after their shortcomings are

exposed (Hacking, 1982). We extend this body of work by showing that accountability metrics remain malleable even after they become institutionalized. The significance of a metric is not exhausted by its operationalization. Accountability is also shaped by the incentive structures that attach consequences to numbers and by the auditing practices through which numbers are produced and interpreted. Disentangling these dimensions helps explain how the same metric (in name) can support competing political projects, justify divergent interventions, and continually redistribute the burdens of accountability across different actors. More broadly, our framework provides a vocabulary for analyzing how accountability systems evolve even when the metrics at their core remain stable. In many ways the SNAP QC system has remained remarkably consistent over its 50 year history. Yet, changes to the surrounding infrastructure have led to different understandings of who is accountable and for what.

Our findings also carry implications for the growing use of algorithmic systems in public administration. Contemporary fraud detection and risk-scoring systems are often presented as new accountability technologies capable of rooting out waste, fraud, and abuse. Yet these systems inherit many of the assumptions and accountability relations embedded in pre-existing quantification infrastructures. Algorithmic systems inherit existing operationalizations of error, reinforce specific incentive structures, and frequently expand auditing practices through intensified of data collection and surveillance. As a result, debates about algorithmic accountability cannot be limited to questions of model performance, transparency, or explainability. They must also attend to the prior choices that determine what counts as an error, whose behavior is scrutinized, which outcomes are prioritized, and who bears the costs of accountability. In the context of public benefits administration, the danger is that algorithmic fraud detection systems may further entrench a conception of accountability grounded in suspicion, surveillance, and payment accuracy while obscuring questions of access, service quality, and governmental responsibility.

Beyond the SNAP case, this paper suggests that accountability metrics should be understood as sites of ongoing political contestation rather than neutral administrative tools. As governments turn to algorithmic systems in public administration, the choices embedded in accountability infrastructures are particularly salient. Examining how metrics are operationalized, incentivized, and audited provides a way to make these choices visible and contestable. Such scrutiny is essential if accountability systems are to address not only administrative error, but also broader democratic commitments to public service.

CHAPTER 6

Conclusion

In this dissertation, I introduce the concept of *legitimacy practices*. I show how everyday administrative practices around agency technology transform what is seen as a legitimate policy intervention, who is seen as a legitimate policy actor, and on what grounds. As the role of agency officials evolves through algorithmic adoption, I examine how the administrative state negotiates its own legitimacy through mundane activities like oversight and management, design and implementation, or quality control. These practices reconfigure the relationships between different policy actors and shape the boundaries of "good" policy.

Through three distinct settings, I show that thorny social and political problems get suppressed and reshaped through technology and the rhetoric that accompanies it. In Chapter 3, I outline agencies' performance of everyday legitimacy practices and describe how the adoption of algorithmic technology in government has shifted these practices to emphasize efficiency and narrow notions of non-arbitrariness rather than accountability. In Chapter 4, I describe how technologies can play a rhetorical role in establishing legitimacy by narrowing the scope of accountability. Finally, in Chapter 5, I examine how accountability is embedded through infrastructural practices, even when accountability has otherwise been sidelined in agency legitimacy practices. The invisibility of accountability infrastructures allow changes (and even perversions) of accountability to occur largely unnoticed.

Taken together, these three studies highlight the political work embedded within algorithmic technologies, contextualizing this work within a longer lineage of the administrative state's battle to establish its own legitimacy. I show that notions of what makes the administrative state legitimate change over time with their practices. By attending to different moments in the history of algorithmic technology, I argue that this isn't a one-off, static change brought about by the adoption of a new technology, but something that evolves over time. This dissertation argues that legitimacy, how it is performed, and by whom, has important implications for the allocation of power within our political processes and the values that sustain the continued existence of the administrative state.

6.1 Reflections and Future Directions

Questions of administrative legitimacy under algorithmic governance are not merely of academic interest. The pressures that agencies face to demonstrate their legitimacy and how they respond to these pressures meaningfully shapes the work that they do, with important consequences for the government's relationship to the public: how well it serves the public interest, how responsive it is to its constituents, and the burdens it imposes on individuals.

Since beginning this dissertation in 2023, the shape of administrative legitimacy has been distorted by constant efforts to undermine agency authority. Algorithms, including AI, have played a significant role in this shift. The growing displacement of human judgment by algorithms is occurring at the precise moment when traditional sources of legitimacy – democratic accountability, procedural fairness, and expert authority– are under sustained political attack. While the Department of Government Efficiency—which introduced AI with the goal of eliminating government waste—was itself short-lived, its implications reflect a longer-term shift in legitimacy under algorithmic governance. In this vision of legitimacy, efficiency becomes its own legitimating principle. Claims of efficiency provide political cover for accountability gaps, hiding the work of political decision-making and sidestepping the messy work of democratic deliberation. Accountability persists through existing government infrastructures, but this new accountability is accountability in name only. Documentation requirements and error metrics serve as accountability processes with no forum passing judgment on the other end.

This accountability to no one explains how algorithmic governance is gaining institutional momentum despite a lack of popular legitimacy. This dissertation provides evidence of a growing gap between legitimacy practices and both empirical acceptance of authority and normative conceptions legitimacy. This marks a shift from democratic legitimacy toward a top-down, managerial conception of legitimacy. Legitimacy can no longer be studied only through relationship between public and state, but must be studied through demonstrations of legitimacy for the state. The work in this dissertation outlines how the government institutions perform legitimacy for other parts of government. In future work, I hope to expand this analysis to examine how non-governmental actors achieve legitimacy in the eyes of the state, for example in the case of corporate involvement in AI evaluation and standards-making and the case of professional associations like the Association for Computing Machinery's participation in technology policy. I also hope to expand the work in Chapter 5 to trace the relationship between metrics and fraud-detection algorithms. Through a public records request, my collaborators and I have access to Fraud Framework Implementation Grant proposals, which will serve as a source of data for understanding the claims that algorithms

advance in addressing government waste. Ultimately, this line of work seeks to clarify what forms of authority become possible when efficiency itself is treated as a sufficient basis for legitimacy. As legitimacy practices continue to evolve alongside administrative technologies, understanding how legitimacy is produced and transformed is essential to understanding what kind of state algorithmic governance is helping to build.

BIBLIOGRAPHY

Amina Abdu and Abigail Jacobs. Public administration. In Jenna Burrell, Ranjit Singh, and Patrick Davison, editors, *Keywords of the Datafied State*, pages 78–92. Data & Society, 2024.

Amina A Abdu, Lauren M Chambers, Deirdre K Mulligan, and Abigail Z Jacobs. Algorithmic transparency and participation through the handoff lens: Lessons learned from the US Census Bureau's adoption of differential privacy. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 1150–1162, 2024.

John M Abowd. Staring-down the database reconstruction theorem. In *Joint Statistical Meetings, Vancouver, BC*, volume 234, 2018.

John M Abowd, Tamara Adams, Robert Ashmead, David Darais, Sourya Dey, Simson L Garfinkel, Nathan Goldschlag, Daniel Kifer, Philip Leclerc, Ethan Lew, et al. The 2010 census confidentiality protections failed, here's how and why. Technical report, National Bureau of Economic Research, 2023.

Bruce Ackerman. *Social justice in the liberal state*. Yale University Press, 1980.

Ifeoma Ajunwa. Automated governance. *North Carolina Law Review*, 101(2):355, 2023.

Madeleine Akrich. The de-scription of technical objects. *Shaping technology/building society. Studies in sociotechnical change*, pages 205–224, 1992.

Anna Alexandrova. Connecting economic models to the real world: Game theory and the fcc spectrum auctions. *Philosophy of the Social Sciences*, 36(2):173–192, 2006.

Micah Altman, Aloni Cohen, Kobbi Nissim, and Alexandra Wood. What a hybrid legal-technical analysis teaches us about privacy regulation: The case of singling out. *BUJ Sci. & Tech. L.*, 27:1, 2021.

Louise Amoore. *The politics of possibility: Risk and security beyond probability*. Duke University Press, 2013.

Bernadette Atuahene. A theory of stategraft. *NYUL Rev.*, 98:1, 2023.

John Langshaw Austin. *How to do things with words*. Harvard university press, 1975.

Kenneth A Bamberger. Normative canons in the review of administrative policymaking. *Yale LJ*, 118:64, 2008.

Kenneth A Bamberger. Technologies of compliance: Risk and regulation in a digital age. *Tex. L. Rev.*, 88:669, 2009.

Kenneth A Bamberger and Deirdre K Mulligan. Privacy on the books and on the ground. *Stan. L. Rev.*, 63:247, 2010.

Karen Barad. Posthumanist performativity: Toward an understanding of how matter comes to matter. *Signs: Journal of women in culture and society*, 28(3):801–831, 2003.

Rodney Barker. *Political legitimacy and the state*. Oxford University Press, 1990.

Emily Barman. *Caring capitalism: The meaning and measure of social value*. Cambridge University Press, 2016.

Lauren Bauer and Diane Whitmore Schanzenbach. Proposed SNAP cuts would permanently undermine recession readiness and responsiveness. Technical report, The Brookings Institution, 2025.

Burcu Baykurt. Gov-tech as capture: public infrastructures under data capitalism. *Information, Communication & Society*, pages 1–16, 2025.

Jack M. Beermann. Loper Bright and the future of Chevron Deference. *William & Mary Law Review Online*, 65:1, 2024.

Elizabeth Popp Berman. *Thinking like an economist: How efficiency replaced equality in US public policy*. Princeton University Press, 2022.

Elizabeth Popp Berman and Daniel Hirschman. The sociology of quantification: Where are we now? *American Sociological Review*, 2018.

Emily Berman. A government of laws and not of machines. *BU L. Rev.*, 98:1277, 2018.

Wiebe E Bijker, Roland Bal, and Ruud Hendriks. *The paradox of scientific authority: The role of scientific advice in democracies*. MIT Press, 2009.

Abeba Birhane. Algorithmic injustice: a relational ethics approach. *Patterns*, 2(2), 2021.

Ed Bolen. A poor measure of the wrong thing: The food stamp program's quality control system discourages participation by working families. *Hastings LJ*, 53:213, 2001.

Emily Adlin Bosk. What counts? quantification, worker judgment, and divergence in child welfare decision making. In *Human Service Organizations and the Question of Impact*, pages 93–112. Routledge, 2020.

Mark Bovens. Analysing and assessing accountability: A conceptual framework. *European Law Journal*, 13(4):447–468, 2007.

danah boyd and Jayshree Sarathy. Differential perspectives: Epistemic disconnects surrounding the us census bureau's use of differential privacy. *Harvard Data Science Review*, 2022.

Lisa Schultz Bressman. Beyond accountability: Arbitrariness and legitimacy in the administrative state. *NYU L. Rev.*, 78:461, 2003.

Vannevar Bush. Science–the endless frontier: A report to the president on a program for postwar scientific research (75th anniversary edition). *Washington DC, National Science Foundation*, 1945/2020.

Madalina Busuioc. Accountable artificial intelligence: Holding algorithms to account. *Public administration review*, 81(5):825–836, 2021.

Ryan Calo and Danielle Keats Citron. The automated administrative state: A crisis of legitimacy. *Emory LJ*, 70:797, 2020.

CBPP. Food stamp overpayment error rate hits a record low. Technical report, Center on Budget and Policy Priorities, 2003.

CBPP. Policy Basics: The Supplemental Nutrition Assistance Program (SNAP). Technical report, Center on Budget and Policy Priorities, 2024.

Brian J. Chen. Dispelling myths of AI and efficiency. *Data & Society policy brief*, 2025.

Daryl E Chubin. Open science and closed science: Tradeoffs in a democracy. *Science, Technology, & Human Values*, 10(2):73–80, 1985.

Hyunsik Chun and Michael Sauder. The logic of quantification: Institutionalizing numerical thinking. *Theory and Society*, 51(2):335–370, 2022.

Dorothy E Chunn and Shelley AM Gavigan. Welfare law, welfare fraud, and the moral regulation of the 'never deserving' poor. *Social & legal studies*, 13(2):219–243, 2004.

Danielle Keats Citron. Technological due process. *Wash. U. L. Rev.*, 85:1249, 2007.

Adele E Clarke, Carrie Friese, and Rachel S Washburn. *Situational analysis: Grounded theory after the interpretive turn*. Sage publications, 2017.

Cary Coglianese and David Lehr. Transparency and algorithmic governance. *Administrative Law Review*, 71(1):1–56, 2019.

Julie E Cohen. Oligarchy, state, and cryptopia. *Fordham Law Review*, 94, 2025.

Harry M Collins. Stages in the empirical programme of relativism, 1981.

A Feder Cooper and Ellen Abrams. Emergent unfairness in algorithmic fairness-accuracy trade-off research. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 46–54, 2021.

A Feder Cooper, Jonathan Frankle, and Christopher De Sa. Non-determinism and the lawlessness of machine learning code. In *Proceedings of the 2022 Symposium on Computer Science and Law*, pages 1–8, 2022a.

A Feder Cooper, Emanuel Moss, Benjamin Laufer, and Helen Nissenbaum. Accountability in an algorithmic society: relationality, responsibility, and robustness in machine learning. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 864–876, 2022b.

Richard J Cox and David A Wallace. *Archives and the public good: accountability and records in modern society*. Bloomsbury Publishing USA, 2002.

Ronald Cramer, Ivan Bjerre Damgård, et al. *Secure multiparty computation*. Cambridge University Press, 2015.

Rachel Cummings, Varun Gupta, Dhamma Kimpara, and Jamie Morgenstern. On the compatibility of privacy and fairness. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*, pages 309–315, 2019.

John Danaher. The threat of algocracy: Reality, resistance and accommodation. *Philosophy & technology*, 29(3):245–268, 2016.

Paul Dans and Steven Groves, editors. *Mandate for Leadership: The Conservative Promise*, 2023. The Heritage Foundation.

Lorraine Daston and Peter Galison. *Objectivity*. Princeton University Press, 2007.

Samuel DeCanio. Efficiency, legitimacy, and the administrative state. *Social Philosophy and Policy*, 38(1):198–219, 2021.

David Demortain. The politics of calculation. towards a sociology of quantification in governance. *Revue d'anthropologie des connaissances*, 13(13-4), 2019.

Lina Dencik, Arne Hintz, and Jonathan Cable. Towards data justice? the ambiguity of anti-surveillance resistance in political activism. *Big Data & Society*, 3(2):2053951716679678, 2016.

William Deringer. *Calculated values: Finance, politics, and the quantitative age*. Harvard University Press, 2018.

Damien Desfontaines and Balázs Pejó. SoK: differential privacies. *arXiv preprint arXiv:1906.01337*, 2019.

Emmanuel Didier. *America by the numbers: Quantification, democracy, and the birth of national statistics*. MIT Press, 2020.

U.S. Digital Service. Mission. https://web.archive.org/web/20250410112907/https://www.usds.gov/mission accessed April 21, 2025.

Bibiana Duarte-Abadía, Rutgerd Boelens, and Emerson Buitrago. Neoliberal commensuration and new enclosures of the commons: mining and market–environmentalism governmentalities. *Territory, Politics, Governance*, 11(7):1480–1500, 2023.

Cynthia Dwork and Moni Naor. On the difficulties of disclosure prevention in statistical databases or the case for differential privacy. *Journal of Privacy and Confidentiality*, 2(1), 2010.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284. Springer, 2006.

Lauren B Edelman. *Working law: Courts, corporations, and symbolic civil rights*. University of Chicago Press, 2020.

Paul N Edwards. *The closed world: Computers and the politics of discourse in Cold War America*. MIT Press, 1996.

David Freeman Engstrom, Daniel E Ho, Catherine M Sharkey, and Mariano-Florentino Cuéllar. Government by algorithm: Artificial intelligence in federal administrative agencies. 2020.

Paul Erickson. *The world the game theorists made*. University of Chicago Press, 2019.

Paul Erickson, Judy L. Klein, Lorraine Daston, Rebecca Lemov, Thomas Sturm, and Michael D. Gordin. *How Reason Almost Lost Its Mind*. University of Chicago Press, Chicago, 2013.

Wendy Nelson Espeland and Michael Sauder. Rankings and reactivity: How public measures recreate social worlds. *American Journal of Sociology*, 113(1):1–40, 2007.

Wendy Nelson Espeland and Mitchell L Stevens. Commensuration as a social process. *Annual review of sociology*, 24(1):313–343, 1998a.

Wendy Nelson Espeland and Mitchell L. Stevens. Commensuration as a Social Process. *Annual Review of Sociology*, 24(1):313–343, 1998b. doi: 10.1146/annurev.soc.24.1.313. URL https://doi.org/10.1146/annurev.soc.24.1.313.

Wendy Nelson Espeland and Mitchell L Stevens. A sociology of quantification. *European Journal of Sociology/Archives européennes de sociologie*, 49(3):401–436, 2008.

Wendy Nelson Espeland and Berit Irene Vannebo. Accountability, quantification, and law. *Annu. Rev. Law Soc. Sci.*, 3:21–43, 2007.

Virginia Eubanks. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press, 2018.

Henry Farrell. AI as governance. *Annual Review of Political Science*, 28(1):375–392, 2025.

Henry Farrell and Abraham L. Newman. Elon Musk Weaponizes the Government. *LawFare*, 2025.

Patricia Moloney Figliola and Jill C Gallagher. The Federal Communications Commission's spectrum auction authority: History and options for reinstatement. *Congressional Research Service (CRS) Reports and Issue Briefs*, pages NA–NA, 2023.

Jill E. Fisch. Overseeing the administrative state. *Seattle University Law Review*, 47:899, 2024.

U.S. Food and Nutrition Service. SNAP Quality Control. https://www.fns.usda.gov/snap/qc, accessed September 2, 2025.

Jeremy Freese and David Peterson. The emergence of statistical objectivity: Changing ideas of epistemic vice and virtue in science. *Sociological Theory*, 36(3):289–313, 2018.

Walter Bryce Gallie. Essentially contested concepts. In *Proceedings of the Aristotelian Society*, 1956.

GAO. Disseminating information on successful use of data analytics could help states manage fraud risks. Technical report, U.S. Government Accountability Office, 2018.

GAO. About. https://web.archive.org/web/20250421191630/https://www.gao.gov/about, 2025. archived and accessed from https://www.gao.gov/about on April 21, 2025.

Michele E Gilman. The class differential in privacy law. *Brooklyn Law Review*, 77(4), 2012.

Jake Goldenfein, Deirdre K Mulligan, Helen Nissenbaum, and Wendy Ju. Through the handoff lens: competing visions of autonomous futures. *Berkeley Tech. LJ*, 35:835, 2020.

Ben Green. Escaping the impossibility of fairness: From formal to substantive algorithmic fairness. *Philosophy & Technology*, 35(4):90, 2022.

Ben Green and Salomé Viljoen. Algorithmic realism: expanding the boundaries of algorithmic thought. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 19–31, 2020.

Emily Gutierrez. Changes to SNAP and Medicaid would have implications for student access to school meals. Technical report, Urban Intsitute, 2025.

Ian Hacking. Biopower and the avalanche of printed numbers. *Humanities in Society*, 5(3-4): 279–295, 1982.

Michael Hammond. What is an affordance and can it help us understand the use of ICT in education? *Education and Information Technologies*, 15:205–217, 2010.

Spencer Headworth. *Policing Welfare: Punitive Adversarialism in Public Assistance*. University of Chicago Press, May 2021. ISBN 978-0-226-77922-5. doi: 10.7208/chicago/9780226779539.001.0001. URL https://doi.org/10.7208/chicago/9780226779539.001.0001.

Spencer Headworth. Stategraft in public assistance programs. *Wis. L. REv.*, page 503, 2024.

Ursula Hébert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. Multicalibration: Calibration for the (computationally-identifiable) masses. In *International Conference on Machine Learning*, pages 1939–1948. PMLR, 2018.

Mireille Hildebrandt. Closure: On Ethics, Code, and Law. In *Law for Computer Scientists and Other Folk*. Oxford University Press, 04 2020. ISBN 9780198860877. doi: 10.1093/oso/9780198860877.003.0011.

Mireille Hildebrandt. *Machine We Trust. Perspectives on Dependable AI*, chapter The Issue of Bias. The Framing Powers of Machine Learning. MIT Press, 2021.

Sun-ha Hong. Prediction as extraction of discretion. *Big Data & Society*, 10(1): 20539517231171053, 2023.

Sarah E Igo. *The averaged American: Surveys, citizens, and the making of a mass public*. Harvard University Press, 2007.

Andrew Iliadis and Amelia Acker. The Palantir Files: Public interest archives for platform accountability. *Information, Communication & Society*, 27(13):2343–2365, 2024.

Abigail Z Jacobs and Deirdre K Mulligan. The hidden governance in AI. *The Regulatory Review*, 2022.

Ron S Jarmin, John M Abowd, Robert Ashmead, Ryan Cumings-Menon, Nathan Goldschlag, Michael B Hawes, Sallie Ann Keller, Daniel Kifer, Philip Leclerc, Jerome P Reiter, et al. An in-depth examination of requirements for disclosure risk assessment. *Proceedings of the National Academy of Sciences of the United States of America*, 120(43):e2220558120, 2023.

Sheila Jasanoff. *Risk management and political culture*, volume 12. Russell Sage Foundation, 1986.

Carlee Joe-Wong, Soumya Sen, Tian Lan, and Mung Chiang. Multiresource allocation: Fairness–efficiency tradeoffs in a unifying framework. *IEEE/ACM Transactions on Networking*, 21(6):1785–1798, 2013.

Lauren Kallins and Ashton Thompson. How states are responding to new SNAP requirements. *National Conference of State Legislatures*, Mar 2026. URL https://www.ncsl.org/state-legislatures-news/details/how-states-are-responding-to-new-snap-requireme

Margot E Kaminski. Understanding transparency in algorithmic accountability. In Woodrow Barfield, editor, *Cambridge Handbook of the Law of Algorithms*, pages 20–34. Cambridge University Press, 2020.

Sallie Ann Keller and John M Abowd. Database reconstruction does compromise confidentiality. *Proceedings of the National Academy of Sciences*, 120(12):e2300976120, 2023.

Makena Kelly. Elon Musk Ally Tells Staff 'AI-First' Is the Future of Key Government Agency. https://www.wired.com/story/elon-musk-lieutenant-gsa-ai-agency/, 2025.

Joshua B Kennedy. Always watchful: Political context and congressional oversight through the Government Accountability Office. *Congress & the Presidency*, 50(3):342–360, 2023.

Christopher T Kenny, Shiro Kuriwaki, Cory McCartan, Evan TR Rosenman, Tyler Simko, and Kosuke Imai. The use of differential privacy for census data and its impact on redistricting: The case of the 2020 US Census. *Science Advances*, 7(41):eabk3283, 2021.

Christopher T Kenny, Shiro Kuriwaki, Cory McCartan, Evan TR Rosenman, Tyler Simko, and Kosuke Imai. Comment: The essential role of policy evaluation for the 2020 census disclosure avoidance system. *Harvard Data Science Review*, 2022.

Daniel Kifer, John M Abowd, Robert Ashmead, Ryan Cumings-Menon, Philip Leclerc, Ashwin Machanavajjhala, William Sexton, and Pavel Zhuravlev. Bayesian and frequentist semantics for common variations of differential privacy: Applications to the 2020 census. *arXiv preprint arXiv:2209.03310*, 2022.

Rob Kitchin. Urban data power: Capitalism, governance, ethics, and justice. In Ola Söderström and Ayona Datta, editors, *Data power in action*, pages 21–41. Bristol University Press, 2023.

Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Cass R Sunstein. Discrimination in the age of algorithms. *Journal of Legal Analysis*, 10:113–174, 2018.

Daniel Lee Kleinman and Mark Solovey. Hot science/Cold War: The National Science Foundation after World War II. *Radical History Review*, 1995(63):111–139, 1995.

Nancy Kranich. Countering enclosure: reclaiming the knowledge commons. In *Understanding knowledge as a commons: from theory to practice*, pages 85–122. MIT Press, 2006.

Joshua A Kroll. Outlining traceability: A principle for operationalizing accountability in computing systems. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 758–771, 2021.

Joshua A Kroll, Solon Barocas, Edward W Felten, Joel R Reidenberg, David G Robinson, and Harlan Yu. Accountable algorithms. *U. Pa. L. Rev.*, 165:633, 2017.

Sapna Kumar. Scientific and technical expertise after Loper Bright. *Duke Law Journal*, 74, 2025.

Sarah Esther Lageson. *Digital Punishment: Privacy, Stigma, and the Harms of Data-Driven Criminal Justice*. Oxford University Press, New York, NY, 2020. ISBN 978-0-19-087200-7.

Paul J. Larkin. Agency deference after *Kisor v. Wilkie*. *Georgeotown Journal of Law & Public Policy*, 18:105, 2020.

Brendan Lawson and Matthew Lovatt. Towards a rhetorical understanding of statistics in politics: Quantifying the national health service 'winter crisis'. *European Journal of Communication*, 36(2):110–124, 2021.

Seth Lazar. Legitimacy, authority, and democratic. *Oxford Studies in Political Philosophy*, 10:28, 2024.

Robert Leonard. *Von Neumann, Morgenstern, and the creation of game theory: From chess to social science, 1900–1960*. Cambridge University Press, 2010.

Karen Levy, Kyla E Chasalow, and Sarah Riley. Algorithms and decision-making in the public sector. *Annual Review of Law and Social Science*, 17:309–334, 2021.

Jenny M. Lewis. The politics and consequences of performance measurement. *Policy and Society*, 34(1):1–12, 2015.

John A Lynch. The rhetoric of technology as a rhetorical technology. *Poroi*, 9(1), 2013.

Mary Madden, Michele Gilman, Karen Levy, and Alice Marwick. Privacy, poverty, and big data: A matrix of vulnerabilities for poor Americans. *Wash. UL Rev.*, 95:53, 2017.

Kirsten Martin and Ari Waldman. Are algorithmic decisions legitimate? the effect of process and outcomes on perceptions of legitimacy of AI decisions. *Journal of Business Ethics*, 183(3):653–670, 2023.

Lawrence L Martin. The rush to measure performance. *J. Soc. & Soc. Welfare*, 25:63, 1998.

Andrea Mennicken and Wendy Nelson Espeland. What's new with numbers? Sociological approaches to the study of quantification. *Annual Review of Sociology*, 45(1):223–245, 2019.

Sally Engle Merry. *The Seductions of Quantification: Measuring Human Rights, Gender Violence, and Sex Trafficking*. University of Chicago Press, 2016.

Sally Engle Merry, Kevin E Davis, and Benedict Kingsbury. *The quiet power of indicators: Measuring governance, corruption, and rule of law*. Cambridge University Press, 2015.

Jacob Metcalf, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, and Madeleine Clare Elish. Algorithmic impact assessments and accountability: The co-construction of impacts. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 735–746, 2021.

Shira Mitchell, Eric Potash, Solon Barocas, Alexander D'Amour, and Kristian Lum. Algorithmic fairness: Choices, assumptions, and definitions. *Annual Review of Statistics and its Application*, 8(1):141–163, 2021.

Donald P. Moynihan, Sergio Fernandez, Soonhee Kim, Kelly M. LeRoux, Suzanne J. Piotrowski, Bradley E. Wright, and Kaifeng Yang. Performance Regimes Amidst Governance Complexity. *Journal of Public Administration Research and Theory*, 21(suppl_1):

i141–i155, January 2011. ISSN 1053-1858. doi: 10.1093/jopart/muq059. URL https://doi.org/10.1093/jopart/muq059.

Deirdre K Mulligan and Kenneth A Bamberger. Saving governance-by-design. *California Law Review*, 106(3):697–784, 2018.

Deirdre K Mulligan and Kenneth A Bamberger. Procurement as policy: Administrative process for machine learning. *Berkeley Tech. LJ*, 34:773, 2019.

Deirdre K Mulligan and Helen Nissenbaum. The concept of handoff as a model for ethical analysis and design. *The Oxford handbook of ethics of AI*, 1(1):233, 2020.

Deirdre K Mulligan, Colin Koopman, and Nick Doty. Privacy is an essentially contested concept: a multi-dimensional analytic for mapping privacy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2083):20160118, 2016.

Deirdre K Mulligan, Joshua A Kroll, Nitin Kohli, and Richmond Y Wong. This thing called fairness: Disciplinary confusion realizing a value in technology. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–36, 2019.

Priyanka Nanayakkara and Jessica Hullman. What's driving conflicts around differential privacy for the U.S. Census. *IEEE Security & Privacy*, 21(5):33–42, 2023. doi: 10.1109/MSEC.2022.3202793.

Dorothy Nelkin. Threats and promises: Negotiating the control of research. *Daedalus*, pages 191–209, 1978.

Edward Nik-Khah and Philip Mirowski. The ghosts of Hayek in orthodox microeconomics: Markets as information processors. In *Markets*. meson press, 2019.

Helen Nissenbaum. Accountability in a computerized society. *Science and Engineering Ethics*, 2:25–42, 1996.

Donald A Norman. Affordance, conventions, and design. *Interactions*, 6(3):38–43, 1999.

Paul Ohm. Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA L. Rev.*, 57:1701, 2009.

Oklahoma Senate. Senator bergstrom files senate bill 1310 to address snap error rate. Technical report, Oklahoma Senate, 2025.

Frank Pasquale. *The black box society: The secret algorithms that control money and information*. Harvard University Press, 2015.

Jule Pattinson-Gordon. Can the right tech and training lower Mississippi's SNAP error rates? *Governing*, Dec 2025. URL https://www.governing.com/policy/can-the-right-tech-and-training-lower-mississippis-snap-error-rates

Fabienne Peter. Rawls' idea of public reason and democratic legitimacy. *Politics and Ethics Review*, 3(1):129–143, 2007.

Trevor J Pinch and Wiebe E Bijker. The social construction of facts and artefacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social studies of science*, 14(3):399–441, 1984.

Theodore M Porter. Quantification and the accounting ideal in science. *Social studies of science*, 22(4):633–651, 1992.

Theodore M Porter. *Trust in numbers: The pursuit of objectivity in science and public life*. Princeton University Press, 1995.

S. James Press and Judith M. Tanur. The Confluence of Sociology, Statistics, and Public Policy in the Quality Control of the Food Stamps, AFDC, and Medicaid Family Assistance Programs. *Evaluation Review*, 15(3):315–332, June 1991. ISSN 0193-841X, 1552-3926. doi: 10.1177/0193841X9101500302. URL https://journals.sagepub.com/doi/10.1177/0193841X9101500302.

RAND. A brief history of the rand corporation. https://www.rand.org/about/history.html, accessed January 20, 2025.

John Rawls. Political liberalism. In *The New Social Theory Reader*, pages 123–128. Routledge, 2020a.

John Rawls. *A theory of justice: Revised edition*. Harvard university press, 2020b.

Ronald Reagan. Remarks Announcing the Intention To Nominate Charles A. Bowsher To Be Comptroller General of the United States, July 1981. https://www.reaganlibrary.gov/archives/speech/remarks-announcing-intention-nominate-charles-bowsl accessed April 23, 2025.

Robert F Rich. Politics, public policy-making, adn the process of reaching closure. In Hugo Tristram Engelhardt and Arthur L Caplan, editors, *Scientific controversies: Case studies in the resolution and closure of disputes in science and technology*, pages 151–167. Cambridge University Press, 1987.

Gernot Rieder and Judith Simon. Datatrust: Or, the political quest for numerical evidence and the epistemologies of Big Data. *Big Data & Society*, 3(1):2053951716649398, 2016.

Joy Rohde. The contest between information and uncertainty: The Cold War origins of computational policy knowledge. *Public Culture*, 35(3):405–416, 2023.

Nicole Santa Cruz. "the alarm bell": Arizona's drop in SNAP participation signals potential nationwide impact of Trump legislation. *Governing*, Apr 2026. URL https://www.propublica.org/article/arizona-snap-benefits-trump-legislation.

Kurt M Saunders. Law as rhetoric, rhetoric as argument. *Journal of Legal Education*, 44(4): 566–578, 1994.

Eric W Schoon. Operationalizing legitimacy. *American Sociological Review*, 87(3):478–503, 2022.

James C Scott. *Seeing like a state: How certain schemes to improve the human condition have failed*. Yale university Press, 1998.

Mark Seidenfeld. Syncopated Chevron: Emphasizing reasoned decisionmaking in reviewing agency interpretations of statutes. *Texas Law Review*, 73:83, 1994.

Miriam Seifter. States, agencies, and legitimacy. *Vanderbilt Law Review*, 67:443, 2014.

Jodi L Short. Regulatory managerialism as gaslighting government. *Law and Contemporary Problems*, 86(3), 2023.

Susan Leigh Star and Karen Ruhleder. Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research*, 7(1), 1996.

Jake Iain Stone and Brent Mittelstadt. Legitimate power, illegitimate automation: The problem of ignoring legitimacy in automated decision systems. *ACM Journal on Responsible Computing*, 2024.

Cass R Sunstein. *The cost-benefit state: the future of regulatory protection*. American Bar Association, 2002.

Cass R Sunstein. Cost-benefit analysis and arbitrariness review. *Harv. Envtl. L. Rev.*, 41:1, 2017.

Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang. Privacy loss in Apple's implementation of differential privacy on macOS 10.12. *arXiv preprint arXiv:1709.02753*, 2017.

Iddo Tavory and Stefan Timmermans. *Abductive analysis: Theorizing qualitative research*. University of Chicago Press, 2014.

G.T. Thompson. Holding states accountable: Incentivizing a more efficient and effective SNAP program, 2025.

Stefan Timmermans and Iddo Tavory. *Data analysis in qualitative research: Theorizing with abductive analysis*. University of Chicago Press, 2022.

Tom R. Tyler. *Why People Obey the Law*. Yale University Press, 1990.

U.S. Congress, House Committee on Agriculture. H.R. Rep. No. 95-464, 1977.

U.S. Department of Agriculture Food and Nutrition Service. Fy 2025 snap fraud framework implementation grant awards. Technical report, Food and Nutrition Service, 2025a.

U.S. Department of Agriculture Food and Nutrition Service. Supplemental nutrition assistance program quality control review handbook. Technical report, Food and Nutrition Service, 2025b.

U.S. Government Accountability Office. Artificial intellignce: Agencies have begun implementation but need to complete key requirements. Technical report, U.S. Government Accountability Office, 2023.

U.S. Government Accountability Office. Improper payments: USDA's oversight of the Supplemental Nutrition Assistance Program. Technical report, U.S. Government Accountability Office, 2024a.

U.S. Government Accountability Office. What GAO Does. https://web.archive.org/web/20241103173154/https://www.gao.gov/about/what-gao-does, 2024b. From Internet Archive, November 3, 2024. According to the Internet Archive, this page was last archived with these quotes as of March 21, 2025, and were removed some time on or before March 28, 2025.

Kristen Vaccaro, Karrie Karahalios, Deirdre K Mulligan, Daniel Kluttz, and Tad Hirsch. Contestability in algorithmic systems. In *Companion Publication of the 2019 Conference on Computer Supported Cooperative Work and Social Computing*, pages 523–527, 2019.

Gerald P Vaughan. Spectrum auctions. https://wireless.fcc.gov/releases/JV_speech_4_2_03.pdf, 2003.

Salomé Viljoen, Jake Goldenfein, and Lee McGuigan. Design choices: Mechanism design and platform capitalism. *Big Data & Society*, 8(2):20539517211034312, 2021.

Ari Waldman and Kirsten Martin. Governing algorithmic decisions: The role of decision importance and governance on perceived legitimacy of algorithmic decisions. *Big Data & Society*, 9(1):20539517221100449, 2022.

Ari Ezra Waldman. Power, process, and automated decision-making. *Fordham L. Rev.*, 88: 613, 2019.

Ari Ezra Waldman. Algorithmic legitimacy. *The Cambridge Handbook of Algorithms*, pages 107–120, 2020.

Maranke Wieringa. What to account for when accounting for algorithms: a systematic literature review on algorithmic accountability. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 1–18, 2020.

Langdon Winner. Do artifacts have politics? *Daedalus*, 109(1):121–136, 1980.

Felix T Wu. Defining privacy and utility in data sets. *U. Colo. L. Rev.*, 84:1117, 2013.