



# Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance

Inioluwa Deborah Raji

rajiinio@berkeley.edu

University of California, Berkeley  
Berkeley, CA, USA

Colleen Honigsberg

Stanford University

Stanford, CA, USA

colleenh@law.stanford.edu

Peggy Xu

Stanford University

Stanford, CA, USA

peggyxu@stanford.edu

Daniel Ho

Stanford University

Stanford, CA, USA

dho@law.stanford.edu

## ABSTRACT

Much attention has focused on algorithmic audits and impact assessments to hold developers and users of algorithmic systems accountable. But existing algorithmic accountability policy approaches have neglected the lessons from non-algorithmic domains: notably, the importance of interventions that allow for the effective participation of third parties. Our paper synthesizes lessons from other fields on how to craft effective systems of external oversight for algorithmic deployments. First, we discuss the challenges of third party oversight in the current AI landscape. Second, we survey audit systems across domains – e.g., financial, environmental, and health regulation – and show that the institutional design of such audits are far from monolithic. Finally, we survey the evidence base around these design components and spell out the implications for algorithmic auditing. We conclude that the turn toward audits alone is unlikely to achieve actual algorithmic accountability, and sustained focus on institutional design will be required for meaningful third party involvement.

## CCS CONCEPTS

• Social and professional topics → Governmental regulations.

## KEYWORDS

auditing, policy, algorithms, accountability, society

### ACM Reference Format:

Inioluwa Deborah Raji, Peggy Xu, Colleen Honigsberg, and Daniel Ho. 2022. Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society (AIES'22)*, August 1–3, 2022, Oxford, United Kingdom. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3514094.3534181>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

AIES'22, August 1–3, 2022, Oxford, United Kingdom

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9247-1/22/08...\$15.00

<https://doi.org/10.1145/3514094.3534181>

## 1 INTRODUCTION

Some of the most stunning findings of algorithmic bias and harm have been uncovered through the exceptional efforts of third parties. Regulators, academics, civil society, investigative journalists, specialized law firms and more have already begun to form a network of external actors participating in the informal and formal third party oversight of algorithmic systems. Together, these watchdogs continue to participate in consequential studies to analyze the biased outcomes of models used in domains as varied as content curation [80], online advertising [40], tenant screening [39, 70], child welfare [12], criminal justice [5, 10, 20, 26, 78], hiring [63], immigration [24], education [25] or public health [15, 56]. The impact of this work is undeniable – leading to voluntary moratoriums [30], and several successful lawsuits [75, 79]; nation-wide waves of regulatory action and public outcry [74]; congressional hearings [3]; and the recalls or radical redesigns of the deployed algorithms [12, 64, 65, 88].

However, despite its clear impact, less attention has been paid to the crucial role of third-party oversight in policy proposals that mention algorithmic auditing more generally. Developments such as the General Data Protection Regulation (GDPR) [86], the proposed Algorithmic Accountability Act [48], and several of the state-wide and municipal proposals throughout the U.S. [50] tend to focus instead on *internal audits* commissioned, executed, paid for and controlled by the companies being targeted.

Many experts nevertheless applaud such developments. Kearns and Roth note that “more systematic, ongoing, and legal ways of auditing algorithms are needed” [36]. Kai-Fu Lee even asserted that such internal audits may be “better than regulations,” and pointed to the exemplar of Environmental, Social, and Governance (ESG) reporting and auditing for corporate social responsibility (CSR).<sup>1</sup>

While the sharp rise in ESG investment strategies is indeed impressive, ESG in many ways challenges, not supports, the turn toward the audits of algorithmic and artificial intelligence (AI) branded products. Much like “AI ethics,” the notion of ESG is expansive – encompassing global warming, human rights, labor practices, executive compensation, diversity, equity, and inclusion and more.

<sup>1</sup>Eye on AI Podcast, Aug. 11, 2021, at 15:00 (“It’s impossible for a government to basically evaluate all companies in person so ... audits would make sense ... ESG has generally been successful ... in guiding positive behavior in corporations ... If there are ESG third party watch dogs ... and if shareholders started to care ... that creates a self-monitoring mechanism that could work perhaps even better than regulations.”)

The potential costs of diffuse ESG audits are high. And due to the ambiguity of what is being audited, ESG certifications risk becoming “cheap talk,” rubber stamping practices without in fact promoting social responsibility.<sup>2</sup> A large literature has pointed to these challenges in ESG investing [8, 19, 37, 42, 59, 61, 71]. There is a clear precedent for situations of auditors compromised by selection, payment, or employment by auditees; auditors struggling to secure access; auditors clumsily navigating diffuse performance standards; and auditees that may elect to share or conceal findings otherwise hidden behind a nondisclosure agreement.

The naive move toward AI audits – without considering facets like third-party auditor scope, independence, access, standards, and transparency – risks ignoring these lessons. Like in ESG, conflicts of interest can compromise the quality of third-party audits of algorithmic systems. Auditing the automated decision systems of a large technology platform can rival the complexity of auditing for human rights violations in a large supply chain. The standards and scope of an algorithmic audit may be unclear – leading to assessments that are unfocused and expensive. The mediation of auditor access is difficult, and the identification of audit targets and scope uncertain. Furthermore, audited companies could outright ignore or overlook audit results. Worse yet – as was observed in the AI hiring tool industry [76, 77] – even well-executed algorithmic audit results may get co-opted by audit targets and used to disguise deeper failures. The controversial facial recognition company Clearview AI, for example, claimed to be “100% accurate across all demographic groups according to the ACLU’s facial recognition accuracy methodology,” but failed to address the ethical and legal challenges that accompany its sourcing of face data, which has been subject to numerous legal actions (e.g., in California, Illinois, the European Union) [29].

In sum, not all audits are created equally. Ultimately, much more is required to make algorithmic audits work, including the intentional design of an audit ecosystem enabling the effective participation of third parties. Existing evidence from other regulatory fields suggests that effective target identification mechanisms, independence in selection and compensation, access to data and AI systems, clear and well-scoped audit standards, and post-audit report transparency will be critical to ensure that audits promote, rather than degrade, overall accountability. This article articulates the current challenges for third party oversight (Section 2), surveys the institutional design of audit systems in non-algorithmic domains (Section 3), and synthesizes the evidence base to draw implications for how to design an effective, robust, and scalable AI audit ecosystem (Section 4). We argue that much of what is missing in the policy discussion of algorithmic auditing – and especially third-party algorithmic auditing – are lessons that can be drawn from existing empirical literature on non-algorithmic audits. Due to practical limitations, we confine our analysis mainly to the U.S. context.

<sup>2</sup>Ross and Seymour [71], noting “many companies seem to be doing the minimum necessary to say they obtained assurance on some of their sustainability reporting and many assurance providers are using standards that are so general and discretionary as to provide the consumers of ESG reporting little practical assurance.”

## 2 THE CURRENT U.S. AI AUDIT POLICY LANDSCAPE

We proceed by clarifying some terminology and then assess the current landscape of AI audits, with a focus on the U.S. context. In this article, we use “algorithm”, and related terms such as “automated decision-making systems” interchangeably with “artificial intelligence (AI)” branded products.

### 2.1 Terminology

Audits evaluate performance relative to expected behavior, as part of a broader accountability process [62]. Audits can be qualitative or quantitative, as system expectations can be articulated and assessed in either mode. The scope of an audit does not always allow for a broader view on the downstream impacts of a system or more open ended reflections on the consequences of deployment beyond the direct outcomes of the system itself. Instead, the audit is most focused on the comparison of the reality of the product’s behavior with respect to clearly articulated expectations, standards and claims.

By convention, audits are classified into three types [14]. “First party” audits are conducted by a company of its own products. Many AI ethics teams can be conceived of as fulfilling a first-party audit role. “Second party” audits are performed by a contractual counterparty or an entity hired by that contractual counterparty. Second party audits typically ensure compliance with contract terms. As applied to AI, public sector procurement of AI products may require a second party audit to assess bias, as has been the case for legislative proposals to regulate facial recognition technology. “Third party” audits are conducted by ostensibly independent parties engaged specifically to conduct the audit, typically subject to pre-determined auditing standards. Facebook, for instance, engaged the non-profit organization Business for Social Responsibility to conduct a human rights impact assessment of Facebook’s entrance into Myanmar in 2018 [45]. Another example would be regulators at the U.S. Food and Drug Administration auditing a medical AI tool prior to market entry. We note that while third party oversight by non-profit organizations, journalists, and academics have, as noted above, played critical roles in exposing weaknesses to AI systems, such forms of oversight are not conventionally considered audits per se.

Another distinction is between internal and external oversight. Internal oversight is conducted by stakeholders with a direct employment or contractual relationship with the auditee [66]. External oversight is conducted by third parties with no direct employment or contractual ties to the audit target [64]. Such external parties, which may be public (ie. tied to governmental or other public-interest institutions) or private sector entities, may have no direct access to the audit target and design the audit with different objectives beyond compliance.

We note that the terminology can present uncertainties. Some might label a consultant or contractor hired by a company to conduct an audit to be a “third party” audit, but such an audit may not be considered a form of “external” oversight due to the compensation scheme. Others may use the term “external” and “third party” interchangeably. These uncertainties are precisely the reason to

focus on the specific institutional arrangements of an audit scheme, as we articulate below.

## 2.2 The Outsized Focus on Internal Audits

With these primitives in mind, it is worth noting several major limitations of the current AI audit landscape.

First, many AI policy developments tend to focus on articulating some guidance on what an internal audit should include and potentially enforcing these as expectations for a bare minimum “form” an audit is meant to take. These instructions around audit execution are ultimately attempts at control on audit quality and meaningfulness for accountability, but also reveal the assumption underpinning many of these proposals that *internal* parties are expected to be the ones conducting compliance audits. For instance, the New York City Council mandate for hiring tools specifies the details of a “bias” audit [50]. The Algorithmic Accountability Act of 2022 lists out specific minimum requirements for what is to be included in the required algorithmic impact assessment (AIAs), including a breakdown of evaluation and documentation expectations. Similarly, the Information Commissioner’s Office released a report of “auditability” guidelines for the implementation of impact assessments required as part of the enforcement of the General Data Protection Regulation (GDPR), describing in detail the types of graphs and calculations required to assess notions of accuracy, fairness and explainability [32]. All of these format and content oriented developments work to restrict and define what is to be included in the audit, focusing on concerns that are preferable for internal auditors to operationalize, and that requires a certain level of access or internal co-operation to achieve.

Specifically, the assumption behind calls for Algorithmic Impact Assessments (AIAs) remains that audit activities are executed by internal stakeholders. Most mainstream interpretations of the AIA continue to center its execution as the responsibility of the assessed organization [67, 73]. Even in cases when a participatory approach is considered, the responsibility of consulting with external perspectives when considering impacts to articulate in an AIA is still framed as the responsibility of internal stakeholders to coordinate and interpret these interactions [53]. Thus, AIAs are often practically executed by similar stakeholders as internal audits, with broader regulatory implications and a purpose framed towards prompting and gauging institutional reflexivity and due diligence, in addition to the audit goals of performance assessment with respect to articulated expectations [47]. Unlike most audits, AIAs are not as strictly defined and tend to include more open-ended inquiry into the impacts of a particular algorithmic deployment [55]. The scope of an AIA is thus potentially broader than a specific product feature, and the assessment goes beyond just evaluation measurements, to include more reflective articulations of the trade-offs and decision-making inherent in the release and development of a particular algorithmic system.

Second, the emphasis on internal audits can omit concerns by communities impacted by algorithmic deployments. Internal accountability approaches are developed in reference to a limited set of compliance requirements and self-determined AI principles [66]. This tends to prioritize at most the concerns of regulators or customers, primarily to reduce liability risk, which may not align with core concerns of impacted populations. While helpful in aligning

internal stakeholders, such an approach is not likely to identify or address the main concerns of external stakeholders [54, 66], particularly if that affected population or its representatives are systematically neglected or under-considered as affected non-users. For instance, companies developing algorithmic policing tools naturally prioritize the needs of the police users of their products, even though the community whose safety is impacted by the use of such tools may have different expectations and preferences [51]. In the same vein, regulatory and institutional interpretations of hate speech often differs dramatically from the understanding of those experiencing hate online [43], leading to deployed moderation tools that further marginalize and censor impacted users rather than protect them [16]. Internal auditors can also sometimes direct their attention to focus audits on concerns that are easier for them to measure and address – for example, Twitter [13], Facebook [28] and Hirevue [31] have all been critiqued for promoting a myopic focus on fairness issues while simultaneously diverting attention from other concerns related to the more difficult to formalize challenges of privacy, misinformation and product validity [21].

Thirdly, there is acute potential for conflicts of interest with first or second party audits. For instance, Google researchers on the internal AI team were dismissed and blocked from publishing critiques on the large-scale language models [18, 27], which would later be revealed to be critical to the company’s future product roadmap [23, 85, 87]. Similarly, the warnings from Facebook researchers on addressing mounting bias issues and misinformation campaigns was reported to have been internally suppressed [57]. Even the work of hired consultants can be manipulated when the publication process is mediated by the company itself, as demonstrated by the conflicted reporting on the Hirevue audit by consulting company ORCAA [21]. As a result of this conflict of interest, the ability of internal stakeholders or hired parties to properly scrutinize, honestly describe and effectively alter the algorithmic development process or its outcomes is limited when the incentives to do so directly contradict corporate preferences.

Fourth, even when external auditing is contemplated, the focus is on a narrow set of potential auditors. For instance, the Digital Services Act restricts external access privileges to academic researchers. Similarly, the Algorithmic Accountability Act in 2022, and the Algorithmic Justice and Online Platform Transparency Act of 2021 [49] contemplate obligations on regulated parties with enforcement by the Federal Trade Commission (FTC). Other proposals put the onus on regulatory agencies to serve the primary role of external auditor [36], but regulatory capacity to conduct such audits is currently lacking [22]. In recognition of this capacity gap, the new Algorithmic Accountability Act of 2022 calls for 75 new staff positions for the Federal Trade Commission.

Last, the limited audit capacity has led to many instances of audit failures. Before 2019 and the intervention of the Algorithmic Justice League’s Gender Shades audit [10], the U.S. National Institute of Technology (NIST) failed to measure demographic differences in their regular Facial Recognition Vendor’s Test (FRVT), thus overlooking for many years that “Asian and African American people were up to 100 times more likely to be misidentified than white men” [26]. Similarly, the U.S. Food and Drug Administration has faced challenges in assessing AI-enabled medical devices,

largely deferring to company benchmarks on limited, retrospective data [89].

In sum, while policy attention has been turning toward audits, serious limitations characterize the current AI audit landscape. A future where AI audits are conducted by company employees and consulting companies alone fails to learn the lessons of how third parties have generated robust oversight of AI systems. Such systems are diverse, impacting a range of industries and stakeholders, and it will require more third party participants, representing a variety of communities and making use of a wide set of tools, to concretely assess a wider range of the harms involved for the implicated stakeholders.

### 3 THE INSTITUTIONAL DESIGN OF AUDIT SYSTEMS

Given the limitations of the current AI audit landscape, we now proceed to examine what we can learn from audit systems in other domains. Our goal is to examine other industry successes and failures to consider how to best approach the design of an effective audit ecosystem. We first articulate main differences in how audit systems are designed. We then survey a range of audit systems and illustrate how such differences are manifested.

We break down the institutional design features into five main considerations: (i) **Target Identification & Audit Scope**: The grounds through which the audit subject is identified, and the audit criteria are scoped; (ii) **Auditor Independence**: The degree to which auditors are able to operate without the influence of potentially conflicted stakeholders (e.g., auditor selection, compensation, and rotation requirements); (iii) **Auditor Privileges**: Any special permissions or protections provided to auditors to facilitate the audit process, and protect auditors against retaliation (e.g., access and legal immunity); (iv) **Auditor Professionalization & Conduct Standards**: Articulated rules for audit methodology and auditor qualifications (e.g., accreditation, certification); (v) **Post-audit actions**: Requirements on the auditor, audit target, or other stakeholders to act on audit results (e.g., disclosure).

We now apply this typology to a range of existing audit schemes. We note that the domains are diverse, ranging from financial services, environmental protection, telecommunications, transportation, and healthcare. Some of these audit schemes already intersect with AI products – the National Transportation Safety Board (NTSB) has provided third party analysis of self-driving car crashes; the Food and Drug Administration (FDA) is already approving AI-enabled medical devices; and many current internal audit templates derive from documentation requirements in these other industries. Some AI products, in that sense, are already subject to incumbent audit policies.

We summarize our survey of audit systems in Table 1. Rows present distinct audit systems (for descriptions, see Appendix Table 2) and columns present the five main design features (for details see Appendix Table 1). We group audit systems by whether they are first, second, or third party audits.

We highlight several patterns. First, we note that nearly all audit programs have specific standards and rules that help to scope the audit. Audits to determine whether a good producer meets organic labeling, for instance, are conducted subject to a federal standard [6]. Second, we note that significant variation exists within the third

party audit category, which may be conducted by public agencies (e.g., nursing home inspections) or private parties (e.g., children's toys), which typically also determines whether auditors are compensated by the audited party. Cases in the third party panel are roughly sorted from schemes that are closest to first and second party audits to schemes that are closest to full-blown public inspections. Due to limited regulatory capacity, public inspections appear more likely to be initiated by some complaint or risk-based heuristic, rather than being universally mandated. Third, all audit systems have some mechanism of access to otherwise confidential information: e.g., access to enter a facility to inspect the premises and records or the reporting or access to financial statements. This stands in sharp contrast to third party oversight of AI products, where access has proven challenging. Fourth, accreditation schemes exist for all instances where the auditee selects and compensates the third party auditor. Rotation requirements, which require auditors to rotate across audit targets to ostensibly foster independence, are employed in a minority of cases. Last, in many cases, the audit scheme contemplates post-audit actions, such as recommendations for remediation and disclosure.

What this survey highlights most of all is that there are many varieties of audits. A third party audit where (a) the auditor is selected and compensated by the auditee (b) to conduct an audit of a scope and access determined by the auditee, and (c) with limited post-audit actions may be functionally indistinguishable from first-party audit, and questionably effective in garnering the expected downstream accountability outcomes.

### 4 THE EVIDENCE BASE FOR EFFECTIVE AUDITS

We now examine the evidence base for which design choices may improve the overall effectiveness of an audit. We organize our discussion along the five design categories noted above. For each, we describe the main audit design considerations, present the evidence base from relevant social science and other industries, and develop implications for the AI context.

#### 4.1 Target Identification & Audit Scope

Identifying audit targets and scope is critical to ensuring that audits are applied to relevant parties, and conducted for the right issues.

##### 4.1.1 Evidence Base.

**Audit Selection.** Table 1 shows that a fair number of audit schemes are universal (i.e., they are mandatory for all entities subject to the jurisdiction of the regulator). When audit resources are scarce, however, it may not be possible to select all parties for an audit. In those instances, risk-based or complaint-based selection of audit targets can help to allocate attention where needed the most. Similar resource constraints prompted the development of “syndromic surveillance systems” in public health [35], adverse event reporting systems for pharmaceuticals [66], and financial fraud complaint databases [53], with the goal of providing a real-time signal about potential problems and thus focusing audit efforts. FDA’s “Adverse Event Reporting System,” for example, serves as a kind of post-marketing safety surveillance program for drugs and devices, flagging products with repeated complaints for secondary inspection.

Audit scope & identification			Independence			Privileges	Profession	Post-Audit Actions	
Subject	Selection	Recurring	Selection & payment	Rotation	Auditor	Access	Accreditation	Rec's	Disclosure
<i>First Party</i>									
Software quality	Elective	✓	Auditee		Private	Auditee documentation; on-site visits		✓	
Environmental compliance (EPA self-audit policy)	Elective		Auditee		Private	Auditee disclosure; EPA further requests		✓	✓
<i>Second Party</i>									
Confectionary supply quality	Elective		Auditee		Private	On-site visits		✓	
<i>Third Party</i>									
Post-market telecommunications certification (FCC)	Random		Auditee		Private	Laboratory tests	✓	✓	
Medical Device Single Audit Program (FDA)	Risk-based		Auditee		Private	On-site visits; documents	✓	✓	
Children's products (CPSC)	Universal	✓	Auditee		Private	Laboratory tests	✓	✓	
Organic Labeling (USDA)	Universal	✓	Auditee		Private	On-site visits	✓	✓	
Financial statements (SEC, PCAOB)	Universal	✓	Auditee	✓	Private	On-site visits; financial statements; governance frameworks	✓	✓	
CA greenhouse gas verification	Universal	✓	Auditee	✓	Private	On-site visits; emissions and product data; interviews	✓		✓
Whistleblower disclosure program (FAA)	Complaint		Regulator		Public	Documentation		✓	
Clinical trial audits (FDA)	Complaint		Regulator		Public	On-site visits; patient records		✓	
Post-market drug surveillance (FDA)	Risk-based		Regulator		Public	Post-approval studies; clinical trials	✓	✓	
Nursing homes (HHS)	Universal	✓	Regulator		Public	On-site visits; clinical records; surveys; interviews	✓	✓	

**Figure 1:** This table classifies audit case studies along a range of different design features. “Subject” refers to what is being audited. “Selection” asks how the related harm was discovered and reported (e.g. whether the audit is mandated, voluntary, or complaint-initiated). “Recurring” asks whether there is an ongoing or periodic auditing requirement. “Selection & payment” refers to the entity that pays and selects the auditor. “Rotation” refers to whether auditors are required to rotate off of the auditing process after a certain period of time. “Access” explains how access is granted to auditors. “Accreditation” asks whether the auditing entity must be accredited, certified, or recognized by an official body before it is permitted to conduct the audit. “Rec’s” refers to whether the auditor has the power to make recommendations or changes to the audited entity. “Disclosure” asks whether the contents of the audit must be made available to the public in some form.

Other examples of incident reporting systems abound, including the FDA’s medical device incident database [24], the Federal Aviation Administration’s Aviation Safety Reporting System [3], NIST’s National Vulnerabilities Database [60], and the Consumer Financial Protection Bureau’s complaint system [64]. These third-party reporting systems can help focus auditing resources on issues that matter most [81].

Third-party reporting is most successful when third parties possess high quality information. When an individual’s income is reported to the IRS by third parties like banks and employers, for example, that individual is far more likely to be tax compliant than if they self-reported their income [41, 47]. Third-party reporting has a similar effect on small businesses [2]. Still, consumer complaints have proven to be a valuable signal when harms can be difficult to observe. One analysis of investigated nursing home complaints,

for example, found that the complaints predicted performance at subsequent inspections [71]. Another research study found that the Vaccine Adverse Event Reporting System (VAERS) database is predictive of true safety concerns with newly released vaccines [26].

Complaint-initiated audits, however, also raise challenges. Voluntary reporting can be unreliable and create reporting bias [4]. Residents in low-income and minority communities, for instance, are less likely to complain about street conditions or “nuisance” issues in a 311 system [48]. Similar biases have been found in recent consumer complaints to the Consumer Financial Protection Bureau (CFPB) [65]—although early CFPB complaints, if anything, appeared to be higher from minority neighborhoods [6].

Companies can themselves adopt incident reporting systems to raise user and public awareness of a system’s use in consequential scenarios. Similar to the “How Am I Driving?” sticker on vehicles

for employee driver feedback [72], such prompts could nudge harm reporting for impacted stakeholders. These notices, however, can easily become overwhelming. Some have noted, for example, that the privacy disclosure pop-ups implemented in response to the EU's ePrivacy Directive have increased "consent fatigue" among users without meaningfully improving user privacy [67, 76]. In some cases, users were purposefully inundated with notices designed to induce them to consent by default [55, 57, 61, 70]. These types of informational disclosures can also distract regulators from more pressing problems. One study found, for instance, that restaurant compliance notices were not representative of actual food safety status, and pushed inspectors to spend outsize resources on adjudicating the accuracy of these notices rather than addressing safety hazards directly [36].

***Audit Precision.*** For an audit to successfully measure performance, its scope should be sufficiently delineated. Practitioners and researchers alike have noted the importance of making sure that the scope of an audit isn't too broad or too narrow [42]. If auditors are tasked with vague inspection mandates, their results may be difficult to understand or translate into enforcement actions or changed practices. Take the inconsistency of ESG climate change disclosures: Eccles et al. (2012) found that, despite the SEC's interpretive guidance, there is significant variation between companies and within industries in climate change reporting, with some companies disclosing no information about climate change, others using boilerplate language in their disclosures, and still others including robust quantitative metrics like GHG emissions and energy use [19]. As a result, the authors recommended that industries emphasize precision in determining audit jurisdiction and develop sector-by-sector standards for what metrics should be included in ESG climate reporting. The Financial Accounting Standards Board (FASB) uses materiality to "allow auditors to narrow the audit scope to focus on only the most important items." Some evidence suggests that auditing can help to standardize information for more consistent reporting [30].

#### 4.1.2 Implications for AI Auditing.

***Audit Selection.*** The literature suggests that there may be benefits to establishing a ***national incident reporting system*** to enable third parties to file complaints about algorithmic systems. Nonetheless, considerations for how to adjust for differences in reporting propensities and actively solicit minority perspectives for incident reporting will be a key component of the system's design. Much like how FDA's system has enabled a wide range of techniques to surface issues, the incident reporting system can focus auditing resources on issues that matter most [81]. Regulators can use this information to identify specific priority targets and assess fees or penalties based on audit results. They could then incentivize third party auditors to inspect repeatedly flagged cases in this database through various means, such as RFPs or bug bounties.<sup>3</sup>

The status quo is that the use of algorithms is rarely disclosed in a transparent manner, making it hard to identify audit targets in the first place [84]. As a result, it can be difficult for regulators and the public to identify incidents of algorithmic harm and the ***public***

<sup>3</sup>This could resemble Twitter's algorithmic bias bug bounty [13], except mediated by public actors, and not the audit target itself.

***disclosure of algorithmic products in use*** is necessary to reveal such issues. Current strategies for harm discovery in AI involve ad hoc processes of notice and reporting, limited to situations of harm raised and escalated in public forums such as social media or the press. Third parties have already taken advantage of a range of interventions to identify and prioritize potential harms. At times, public procurement process disclosure requirements such as the Public Oversight of Surveillance Technology (POST) Act [69] will require that certain (often public) institutions such as the New York Police Department (NYPD) publicly release information on which AI technologies are being used. System of Records Notices (SORN), which are intended to notify the public whenever a federal agency creates, modifies, or eliminates records involving personally identifiable information, may alert parties to potential uses of algorithms. Investigative journalism tools such as Algorithm Tips [83], a database of algorithmic tools being used by public agencies, can also make information on algorithmic deployment more publicly accessible. However, even in instances when the algorithmic systems are themselves disclosed, it remains difficult to anticipate which of these deployments are linked to known incidents of harms, and thus challenging to prioritize cases or pursue targets proactively [64]. When communities identify issues with algorithmic deployments, it can be difficult for them to articulate or find concrete evidence for such harms [35]. Even when connected to representative organizations, such individuals or communities often lack the vocabulary and ability to flag the issues they observe [34].

***Audit Precision.*** Audits should be ***scoped precisely***. The Digital Services Act is meant to regulate "recommender systems" developed by very large online platforms (VLOPs). The Algorithmic Accountability Act of 2022, meanwhile, applies to deployed "automated decision systems" (ADS) used in an "augmented critical decision process". Although definitions are provided for these terms, there is little guidance for how regulators, auditors, and companies are meant to translate such descriptions into specific performance expectations. No auditor can audit an entire VLOP single-handedly, and the definition of ADS is so context-dependent and purposefully broad [68] that it does little to help scope an audit. One obvious place to start to narrow the scope is with existing law. Instead of asking, "Is there bias?", an audit may more productively ask, "Does the advertising platform violate disparate treatment provisions of the Fair Housing Act?" [82]. Opening audits up for third party participation could also allow for a variety of stakeholders to come in with diverse yet precise inquiries specific to the context, focus and organizational target of which they are most concerned.

## 4.2 Independence

The key to third-party audits is independence. Minimizing conflicts of interest and undue influence helps to ensure the reliability of audit results.

#### 4.2.1 Evidence Base.

***Cross-selling of non-audit services.*** Research suggests that auditors who may try to sell additional non-auditing services to their clients are more likely to produce client-favorable audits. Enron's payment of audit and consulting fees to Arthur Andersen, coupled with Andersen's fear of losing a very lucrative client, is widely seen

to have contributed to the reckless accounting practices that overstated profits by \$1.5 billion [33, 56]. This led to the development of the Public Company Accounting Oversight Board (PCAOB), created by the Sarbanes–Oxley Act of 2002, to oversee financial auditors for public companies.

Causholli et al. (2014) found that financial auditors whose clients express willingness to purchase their non-audit services in the future tend to produce lower-quality audits [13]. Another study found that private car emissions testing facilities with more opportunities to cross-sell services to customers were more likely to pass cars that did not meet emissions standards [63]. Murphy & Sandino (2009) found that compensation consultants recommended higher CEO salaries when those consultants offered additional services to those executives [59]. Another study, however, concluded that while cross-selling opportunities could lead to lower quality statutory audits, they did not have an impact on the quality of independent expert reports provided to takeover targets—though the incentive to cross-sell could have an impact on lowering fees [9].

***Selection and Compensation.*** There is strong social science demonstrating that selection and compensation can undermine independence and, in turn, audit quality. A two-year randomized controlled trial showed that environmental third party audits were more accurate when auditors were not paid by the firm being audited, but by a common pool of government distributed funding [18]. Similarly, when car owners are given the option to choose their own car inspectors, they select for more lenient car inspectors, resulting in a race to the bottom [10, 34]. One study found that vehicle owners were significantly less likely to keep using car emissions testing stations that fail to report favorably on their vehicles [10, 39, 63]. Similar dynamics have been found in credit rating agencies [20, 40], supply chain audits [69], and accounting audits [58]. In short, auditors should ideally not be selected or paid directly by auditees.

***Auditor tenure.*** The tenure of auditors is another key determinant of independence [15]. There is strong evidence that auditor-client familiarity can induce auditors to make judgments that are favorable to their client. One analysis of nearly 17,000 supplier audits found that supply chain monitors are more lenient when auditing entities they have inspected before [69]. Another study found that credit ratings teams were less accurate in their evaluations when they interacted directly with clients [29]. One approach to reduce concerns regarding familiarity is seen in financial auditing, where independence rules prohibit close relationships between auditors and their clients, and prevent the auditor from assisting in the preparation of the financial statements. If an auditor violates these independence rules, regulators can impose a range of consequences ranging from monetary sanctions to a bar from the industry [1].

Some have proposed mandatory audit rotations to counteract this problem. Sarbanes–Oxley, for example, requires mandatory rotation of the lead audit partner. The evidence on mandatory rotations is more mixed. One study observed relatively few benefits to financial audit partner rotations, though it found that the number of financial restatements issued increased in the first two years after rotation, suggesting some “fresh look” benefits [27]. Other studies focused on the bias-expertise tradeoff [22, 32, 38, 43, 65]. The longer an

auditor gets to know a firm, the more firm-specific expertise the auditor acquires. For example, it might take significant time simply to understand the data ecosystem for housing advertisements on a large platform. This tradeoff has also been documented in food safety inspections, where frequent interactions between inspectors and facilities may improve inspector expertise, but can increase the likelihood of social coercion and even bribery [49, 68].

**4.2.2 *Implications for AI Auditing.*** The AI policy landscape abounds with divergent definitions of “independent audits.” The European Commission’s Digital Services Act proposal defines “independent” auditors as being paid and chosen directly by the audit target. A New York City Council policy proposal requiring that automated employment screening tools undergo “an impartial evaluation by an independent auditor,” fails to even clarify how this auditor is selected and vetted.

Recent reports have found serious limitations with the voluntary audits undertaken by AI vendors in the hiring space, such as Pymetrics and HireVue. In the case of Pymetrics, auditors introduced a framework called the “cooperative audit,” which allowed the audit target—Pymetrics—to help frame what questions the auditors would ask, compromising independence. Similarly, Hirevue commissioned and controlled the scope of its audit by ORCAA, including by mediating any public communication of audit results through their own public relations channels. Voluntary audits like these can also be compromised if AI vendors require pre-publication review of audit reports, or if (like HireVue) they mandate that individuals sign an NDA in order to view the audit results.

Rather than allowing companies to pick their own auditors or dictate the focus of their own audits, the literature suggests that there are significant benefits to a regulator (or third party accreditation body) choosing the auditor or using a central fund, instead of relying on direct payment from audit targets to pay auditors. In particular, the PCAOB experience lends support to the creation of ***an audit oversight board***, housed within a public agency or operating as a separate public interest institution, in order to monitor auditors and to determine the expected scope and standards for audit work. In AI, auditors retained by technology companies are potentially compromised by nondisclosure agreements, financial conflicts, and competition for services. An oversight board could define rules around independence, including whether to (a) mandate auditor rotation, (b) discipline auditors who have violated applicable standards, and (c) prohibit auditors from selling non-audit services to auditees.

### 4.3 Auditor Access

For third-party audits to be effective, auditors must be granted access to enough data to conduct a robust review. This access can mean physical access to certain premises (as in nursing home inspections) or records and documentation (as in financial audits).

**4.3.1 *Evidence Base.*** As Table 1 shows, third party audits invariably involve privileged access to information. Agencies like the Food and Drug Administration, the National Transportation Safety Board, Federal Trade Commission, and the Consumer Financial Protection Bureau commonly investigate cases of safety risks or

performance failures that can harm consumers. Similarly, the Department of Justice and the Department of Housing and Urban Development have engaged in audits to analyze disparate outcomes and compliance to civil rights guidelines for decades [78].

The existing literature has addressed the question of access in a number of ways. Some studies, for example, have examined whether unannounced physical site visits and inspections yield more insight into performance quality than announced inspections, highlighting a concern that audit targets might be able to “game” inspections if they know when they are taking place. One study of 205 NHS hospitals in England found that hospital cleanliness (as perceived by patients) increased during periods when inspections occurred [75]. Another study found that unannounced inspections provided better information about the quality of care at nursing homes, and could reduce the regulatory burdens of conducting inspections—though announced inspections could be better suited to identifying organizational issues [46].

Other research has focused on how differing degrees of audit access influence audit quality. One systematic analysis of the monitoring methods employed by the company PricewaterhouseCoopers (PwC), for example, found that the resulting audits omitted several major violations and contained serious deficiencies [62]. The author attributed these shortcomings in part to management pre-screening of information and interviews for auditors. Another study found that auditors who were subject to PCAOB inspections produced higher quality audits compared to auditors (typically auditors of foreign SEC registrants) that the PCAOB was not allowed to inspect, suggesting that providing access to parties beyond auditors, such as regulators, may provide additional benefits [51]. Some have also proposed increasing auditor access and visibility into a company’s financial records by updating the AICPA’s Generally Accepted Auditing Standards (GAAS) to allow for data analysis during a financial audit [74].

**4.3.2 Implications for AI Auditing.** Currently, third parties are often forced to operate at great legal risk to gain access to the audited AI system, keeping them vulnerable to various forms of corporate retaliation or obstruction. For instance, those making use of crowd-sourcing data collection tools can be interrupted abruptly by claims of violations to privacy or intellectual property rights—NYU researchers studying the Facebook ad delivery algorithm were presented with a cease-and-desist order [58]; German auditors at Algorithm Watch were legally threatened by Instagram [9]; and civil society watch dogs were subpoenaed by Clearview AI [46]. Similarly, several investigators, including academic researcher Christian Sandvig and investigative journalist Julia Angwin, have joined the ACLU to campaign against the use of anti-hacking laws such as the Computer Fraud and Abuse Act (CFAA) to block audit studies that use data scraping [1, 2, 33]. Corporations can also use non-legal strategies to block access, including setting up a paywall, adopting prohibitive Terms of Service conditions, or structuring the product to obscure any clear set of test points. To date, it still often requires extraordinary effort for investigators to gain sufficient access for a thorough analysis. The ACLU—representing the National Fair Housing Alliance (NFHA), Communications Workers of America (CWA), several regional fair housing organizations, and individual consumers and job seekers had to demand that Facebook give them

access to analyze the discrimination in its customized ad delivery system as part of a legal settlement [75].

Even third parties who are contracted to conduct an audit can run into challenges, as access is provided only on the audited target’s terms. For instance, a recent voluntary audit of HireVue, a vendor of algorithmic interviewing and hiring services, was found to be “highly limited”: the audit reviewed HireVue’s documentation of just one of its job candidate assessments, and did not include independent evaluation of HireVue’s data or models. Similarly, attempts by social scientists to work directly with Facebook to gain access to data were thwarted by delayed communications, incomplete data releases and the lack of a clear data dictionary, making the released information indecipherable [60]. As one participant notes, “This was the most frustrating thing I’ve been involved in, in my life.” [44]. Often, such collaborative arrangements also include some control or restriction through contracts of how much of the auditor results can be shared more broadly (as was the case for the Pyometrics audit [88]).

Lack of access to data and algorithmic systems strikes us as the most significant vulnerability of the current AI audit ecosystem. Protecting proprietary information is not a proper response, as all audit systems provide some form of privileged access to auditors, and disclosure does not need to be direct nor absolute. The National Institute of Standards and Technology, for instance, protects models by having companies run models via a custom Application Program Interface (API) for the Face Recognition Vendor Test (FRVT). Such **mediated access**, subject to auditor vetting (perhaps by an audit oversight board) and consistent with audit scope, will be critical to enabling third party auditing of AI systems.

#### 4.4 Professionalization

Training, standardization, and accreditation of auditors are also key to ensuring that audits are accurately and consistently conducted.

**4.4.1 Evidence Base.** Training improves audit quality and more experienced team members are associated with more effective audits [5]. One study of supply chain monitors found that audit teams who had more training found significantly more violations than audit teams that were less well-trained [69]. Another study found that, observing the same conditions, food safety inspectors disagreed on how to implement the health code 60 percent of the time [37]. A randomized controlled trial showed that the accuracy and consistency of these inspections improved with training and peer review. Additional research has found that audit quality improves with the amount of client-specific expertise held by audit team members [5].

Other research shows that clear and rigorous standard-setting can improve the quality of third-party audits. Absent standardization, audit variation can lead audited entities to “shop around” for favored auditors [14]. One study, for example, found that a shift in Belgian financial audit standards towards more rule-based procedures decreased financial auditing errors and increased auditor independence [11]. Another found that PCAOB inspections, which assess the compliance of registered public accounting firms with professional audit rules and standards, improved audit quality by prompting lower-quality auditors to exit the market [17]. However, studies of accounting standards reveal that tightening the strictness of standards could potentially be a double-edged sword—although

effective at aligning auditors with external expectations, overly detailed standards can also lead to complacency by reducing auditor agency and thus promoting a check-boxing “compliance mentality,” which disincentivizes skill development and overall individual competence in outlier situations [25].

As seen in the instances where auditees choose auditors in Table 1, accreditation is widespread. The FDA, for example, may withdraw an auditor’s accreditation if that auditor certifies a supplier that is later linked to an outbreak of foodborne illness or if it refuses to give the FDA access to its records [54]. The literature distinguishes between accreditation, which is typically conducted by a non-governmental body, and licensing administered by the public sector. While much empirical work has focused on the costs and barriers to entry imposed by occupational licensing, less empirical work has probed whether professional licensing has an impact on quality [44]. Existing research on this question is mixed. Several studies have found that increasing or decreasing licensing requirements for CPAs had no meaningful effect on CPA quality [7, 12, 16]. Another analysis of dental licensing requirements found that more rigorous licensing standards did not improve average dental outcomes of patients [45]. One study of occupational licensing laws for teachers found that stricter occupational licensing regimes is associated with a wider distribution of student test scores [52], though others concluded that state occupational licensing requirements for teachers had no impact on students’ standardized test scores [28].

In general, making audit outcomes accessible to the public through open release or a vetted request process improves auditor accountability. For example, some financial auditors must make their communications and records publicly accessible to shareholders in certain U.S. states—this has been proven to improve auditor conduct and lead to higher quality communications [77].

**4.4.2 Implications for AI Auditing.** Given the range of actors who currently provide third party oversight, it will be important to set standards for what counts as an audit and what constitutes acceptable behavior for such a diverse set of potential third party auditors. These auditors make use of a range of tools and methodologies to execute their work, and there is no clear sense of expectations as to what a quality audit entails or who is qualified to execute audit work in the algorithmic context.

The literature strongly supports **training, standardization, and accreditation for third-party AI auditors**. The literature is much more mixed on the benefits of licensing schemes; accreditation through a professional organization may be preferred for three reasons. First, the need for AI audit services is likely to grow significantly in the near future, and substantial evidence suggests that occupational licensing reduces the supply of services. Second, a fully public inspection system seems unlikely in the near term, as the scale, expertise, and agility needed in the AI governance space far outstrips the current capacity of government agencies [22]. Third, accreditation is not exclusive of some regulatory oversight. Recall that Enron inspired a re-assessment of Generally Accepted Auditing Standards (GAAS)<sup>4</sup> and prompted greater oversight of

<sup>4</sup>Note that these are standards to judge and guide auditor practice, determined in the US by the Auditing Standards Board of the American Institute of Certified Public Accountants (AICPA). This differs from Generally Accepted Accounting Principles (GAAP), which are standards to guide corporations in their preparation of financial statements.

auditors by PCAOB. Similar inciting incidents in the medical device industry led to interventions such as the Medical Device Single Audit Program (MDSAP) and to oversight mechanisms for the related Auditing Organizations (AOs) [11].

The capacity gap also strongly calls into question the focus of current policy proposals that focus on academic researchers [60]. Certification should extend to a much wider range of potential auditors, including public interest groups, law firms, and journalists. Limiting audit access to academic researchers may narrowly focus audit resources and attention on questions that are important for academic novelty, but misaligned with the needs of other communities. Non-academic groups can also take advantage of a broader range of strategies to conduct audits and understand harms to distinct communities.

Last, **professional standards** may also determine best practices with respect to legal immunity. Fears of legal repercussions or corporate retaliation can weaken the audit inquiry, and professional standards can help determine limited conditions for liability. Serious violations, such as selling audit target information to competitors, conducting non-audit related activities with the accessed data, or failing to adhere to security or privacy measures, could lead to revocation of credentials.

## 4.5 Post-Audit Actions

After an audit has been completed, investors, regulators, and other affected individuals may want to access and utilize findings to inform shareholder, enforcement, and private actions. We consider here what steps auditors and auditees should be obligated to take after an audit is completed. We focus particularly on the availability and disclosure of audit results.

**4.5.1 Evidence Base.** Public visibility of audit results can improve accountability in several ways. It can prevent companies from being able to disguise or hide undesirable audit outcomes [21] and incentivize better behavior. The Insurance Institute for Highway Safety, for example, operates a vehicle safety database that allows users to search for crash test results for various makes and models. The drug company SuperGen saw its stocks plummet after the FDA revealed that it had overstated the efficacy and safety of its cancer drug. One study found that mandated disclosure of food inspection results increased the compliance of food businesses with food safety regulations [8], though others have found similar food safety disclosure schemes to be ineffective [31, 36].

Separately, the public release of third-party audits can prompt more rigorous auditing practices by allowing for third-party verification. One study found that auditors who have been publicly recognized for their accuracy tend to produce more accurate audits than their non-recognized counterparts, particularly when there are conflicts of interest [23]. An analysis of PwC’s auditing practices also found that the confidential nature of PwC’s audits shielded them from verification by other researchers or NGOs and limited their ability to improve the conditions of the regulated entity [62]. Securities law is premised on the theory that mandated disclosure will improve corporate behavior, and SEC’s EDGAR database allows users to search corporate filings, find out who audited which company, and access audit reports. Recent changes even provide the name of the audit partner.

Mandatory disclosure requirements, however, are not costless. One study found that publicly traded companies spend an average of \$2.2 million to comply with the Sarbanes-Oxley Act's financial reporting requirements [50]. Another analysis of restaurant grading across ten jurisdictions found that the practice—meant to protect consumers and improve food sanitation via targeted disclosures—suffered from several serious deficiencies, including grade inflation and grading inconsistencies. The implementation of a restaurant grading scheme shifted valuable inspection resources away from more significant safety hazards [36]. Other commentators have also noted the tension that exists, particularly in the context of algorithmic accountability, between the desire for AI transparency and the desire to enforce existing intellectual property rights. Broadly scoped AI audits—including things like “ethics audits”—could also be expensive to conduct, particularly for newer startups.

One way of addressing these concerns is to adopt a restricted transparency regime. Sarbanes-Oxley, for example, incentivizes audit firms to fix defects identified in PCAOB inspections by agreeing not to publish deficiencies that are remediated by the firm within 12 months. The FDA reviews applications for medical device pre-market approval (PMA), but does not release information about the devices until after a decision has been rendered. Even then, the FDA releases only a summary about the safety and efficacy of the medical device, along with limited data from the PMA file after all sensitive and confidential information has been removed.

**4.5.2 Implications for AI Auditing.** A restricted transparency regime for registered audit reports could have significant potential benefits. To establish public accountability, all audit results should be issued in a report, and such reports registered and made available in summary form (either publicly or upon vetted request) via a database of audit results. A *registry of audit results* avoids the bias from pre-publication review or the secrecy of findings under NDAs and ensures that audits can provide market incentives for companies to correct behavior, inform enforcement actions, and provide an evidence base for standard setting. Such reports would also help cure one of the major emerging informational asymmetries in technology regulation, by opening the black box of algorithmic design. This would also alert regulators directly of audit outcomes, and open communication channels between them and auditors.

Limited disclosure of audit results would balance the protection of proprietary information against the benefits of external accountability. For such reasons, regulators may already be entitled to fuller access and shareholders of publicly traded companies may inspect audit reports under certain circumstances.

**Repeated assessments**, including after product release, can be effective at ensuring compliance. For instance, it was recurring audits that revealed that Facebook was continuing to discriminate against minority users in job and housing ads, despite earlier communications that it would no longer do so [4, 38, 41, 52, 72].

Audit results can also inform more general **standard-setting** around AI performance and engineering practice. Many AI ethics and accountability discussions have been stymied by the lack of concretely stated expectations of performance and evidence [54]. A major advantage of the audit approach we articulate is its tangibility: audits can help reveal concrete issues and repeated product failures

that can later be resolved by general standards. For instance, IEEE P7013 Inclusion and Application Standards for Automated Facial Analysis Technology industry standard is directly informed by the Gender Shades audit.

## 5 LIMITATIONS

Here we note several important limitations to our work.

First of all, our case studies and evidence offered above are not exhaustive. As just a first step in the direction of learning from cogginate audit schemes, and because the evidence draws from so many different areas, our survey is necessarily selective and incomplete.

Additionally, although our article is critical of internal compliance structures, our work is by no means meant to deprecate the value of internal first-party audits, which is also well understood in other industries such as finance [7]. We only emphasize the role of external and third-party audits in order to illustrate why we cannot rely on internal audits *alone*.

A final comment is that perhaps we do not go far enough in this article. For instance, while we have synthesized the evidence base around each distinct institutional design element, the actual effect may be interactive. Auditor access may only matter when auditors are also independent. We recognize this and believe future work could fruitfully examine the interaction of these dimensions. Also, due to the fluidity of the space, we have purposely stayed away from suggesting whether these reforms should be promulgated via legislation, regulation, enforcement, or self-regulation, although that is an important next step. Our point in this specific text is more simple: as the AI community turns toward audits, it has largely ignored what other audit schemes have to teach us on fostering policy environments conducive to third party participation.

## 6 CONCLUSION

In this article, we demonstrate the dangers of relying on a magic incantation of “audits” to address algorithmic accountability challenges. The mere insistence for audits is not enough – in particular, specific interventions will be necessary to allow for the effective participation of third parties, who play a critical role yet continue to face serious and often debilitating challenges when engaged in their work. AI policy can foster third party audits with more deliberate institutional design, informed by the audit schemes from other industries. Our specific recommendations point to interventions that can foster an ecosystem in which third parties can not only survive in their role, but thrive in directly confronting, verifying, and subjecting to scrutiny the performance claims made by corporations, while adequately addressing complaints of harm brought forth by the impacted population. Our hope is to directly contribute to the growth and maturity of a third party audit ecosystem that can provide the thorough external scrutiny required to protect those currently most vulnerable to harm.

## ACKNOWLEDGMENTS

We thank participants at the 2022 Fall Conference at Stanford's Institute for Human-Centered AI, Fiona Scott Morton, Cathy O'Neil, and DJ Patil for early conversations sparking these ideas, and the Public Interest Technology University Network, the McGovern Foundation, the Algorithmic Justice League, and the Mozilla Foundation for support.

## REFERENCES

- [1] 2020. Sandvig v. Barr. , 73 pages.
- [2] 2020. Van Buren v. United States. , 2667 pages.
- [3] 2021. Impact. <https://themarkup.org/series/impact> Last accessed 25 October 2021.
- [4] Muhammad Ali, Piotr Sapiezynski, Miranda Bogen, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2019. Discrimination through optimization: How Facebook's Ad delivery can lead to biased outcomes. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–30.
- [5] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. *ProPublica*, May 23, 2016 (2016), 139–159.
- [6] Ann H. Baier and Lisa Ahramjian. 2020. Organic Certification of Farms and Businesses Producing Agricultural Products. , 8 pages.
- [7] Jean Bédard and Yves Gendron. 2010. Strengthening the financial reporting system: can audit committees deliver? *International journal of auditing* 14, 2 (2010), 174–210.
- [8] Riccardo Boffo and Robert Patalano. 2020. Esg investing: Practices, progress and challenges. *Éditions OCDE, Paris* (2020).
- [9] Russell Brandom. 2021. Facebook shut down German research on Instagram algorithm, researchers say.
- [10] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [11] Yi-Jung Chen, Chi-Ming Chiu, Yu-Wen Huang, Pei-Weng Tu, Yung-Chuan Lee, and Chia-Hung Chien. 2018. A comparative study of medical device regulations: US, Europe, Canada, and Taiwan. *Therapeutic innovation & regulatory science* 52, 1 (2018), 62–69.
- [12] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*. PMLR, 134–148.
- [13] Rumman Chowdhury and Jutta Williams. 2021. Introducing Twitter's first algorithmic bias bounty challenge. [https://blog.twitter.com/engineering/en\\_us/topics/insights/2021/algorithmic-bias-bounty-challenge](https://blog.twitter.com/engineering/en_us/topics/insights/2021/algorithmic-bias-bounty-challenge) Last accessed 2 November 2021.
- [14] Craig Cochran. 2015. *ISO 9001: 2015 in plain English*. Paton Professional.
- [15] Amanda Coston, Neel Guha, Derek Ouyang, Lisa Lu, Alexandra Chouldechova, and Daniel E Ho. 2021. Leveraging Administrative Data for Bias Audits: Assessing Disparate Coverage with Mobility Data for COVID-19 Policy. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 173–184.
- [16] Ángel Diaz and Laura Hecht. 2021. Double Standards in Social Media Content Moderation. (2021).
- [17] Ilija Djekic, Sladjana Dragojlovic, Zoran Miloradovic, Snezana Miljkovic-Zivanovic, Marija Savic, and Vesna Kekic. 2016. Improving the confectionery industry supply chain through second party audits. *British Food Journal* (2016).
- [18] Christoph Ebelt, Ricardo Baeza-Yates, Richard Benjamins, Hengjin Cai, Mark Coekelbergh, Tania Duarte, Merve Hickok, Aurelie Jacquet, Angela Kim, Joris Krijger, et al. 2021. Towards intellectual freedom in an AI Ethics Global Community. *AI and Ethics* 1, 2 (2021), 131–138.
- [19] Robert G Eccles, Michael P Krzus, Jean Rogers, and George Serafeim. 2012. The need for sector-specific materiality and sustainability reporting standards. *Journal of applied corporate finance* 24, 2 (2012), 65–71.
- [20] Laurel Eckhouse, Kristian Lum, Cynthia Conti-Cook, and Julie Ciccolini. 2019. Layers of bias: A unified approach for understanding problems with risk assessment. *Criminal Justice and Behavior* 46, 2 (2019), 185–209.
- [21] Alex C Engler. 2021. Independent auditors are struggling to hold AI companies accountable. *FastCompany*.
- [22] David Freeman Engstrom and Daniel E Ho. 2020. Algorithmic accountability in the administrative state. *Yale J. on Reg.* 37 (2020), 800.
- [23] William Fedus, Barret Zoph, and Noam Shazeer. 2021. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961* (2021).
- [24] Foxglove. 2020. Home Office says it will abandon its racist visa algorithm – after we sued them. <https://www.foxglove.org.uk/2020/08/04/home-office-says-it-will-abandon-its-racist-visa-algorithm-after-we-sued-them/> Last accessed 25 October 2021.
- [25] Foxglove. 2020. We put a stop to the A Level grading algorithm! <https://www.foxglove.org.uk/2020/08/17/we-put-a-stop-to-the-a-level-grading-algorithm/> Last accessed 25 October 2021.
- [26] Patrick Grother, Mei Ngan, and Kayee Hanaoka. 2019. *Face recognition vendor test (fvrt): Part 3, demographic effects*. National Institute of Standards and Technology.
- [27] Karen Hao. 2020. We read the paper that forced Timnit Gebru out of Google. Here's what it says. *Retrieved January 21 (2020)*, 2021.
- [28] Karen Hao. 2021. How Facebook got addicted to spreading misinformation.
- [29] Caroline Haskins, Ryan Mac, and Logan McDonald. 2020. The ACLU Slammed A Facial Recognition Company That Scraps Photos From Instagram And Facebook. *Buzzfeed News* (2020).
- [30] Rebecca Heilweil. 2020. Big tech companies back away from selling facial recognition to police. That's progress. *Recode*. Available at: <https://www.vox.com/recode/2020/6/10/21287194/amazon-microsoft-ibm-facial-recognition-moratorium-police> (accessed 18 June 2020) (2020).
- [31] HireVue. 2021. HireVue Leads the Industry with Commitment to Transparent and Ethical Use of AI in Hiring.
- [32] ICO. 2020. AI Auditing Framework.
- [33] Thomas Kadri. 2020. Digital Gatekeepers. *Texas Law Review* 99 (2020).
- [34] Michael Katell, Meg Young, Dharma Dailey, Bernease Herman, Vivian Guettler, Aaron Tam, Corinne Bintz, Daniella Raz, and PM Krafft. 2020. Toward situated interventions for algorithmic equity: lessons from the field. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 45–55.
- [35] Michael Katell, Meg Young, Bernease Herman, Dharma Dailey, Aaron Tam, Vivian Guettler, Corinne Bintz, Daniella Raz, and PM Krafft. 2019. An Algorithmic Equity Toolkit for Technology Audits by Community Advocates and Activists. *arXiv preprint arXiv:1912.02943* (2019).
- [36] Michael Kearns and Aaron Roth. 2020. Ethical algorithm design should guide technology regulation. *Retrieved from Brookings*: <https://www.brookings.edu/research/ethical-algorithm-designshould-guide-technology-regulation> (2020).
- [37] Soohun Kim and Aaron Yoon. 2020. Analyzing active managers' commitment to esg: Evidence from united nations principles for responsible investment. *Available at SSRN 3555984* (2020).
- [38] Sara Kingsley, Clara Wang, Alex Mikhalenko, Proteeti Sinha, and Chinmay Kulkarni. 2020. Auditing digital platforms for discrimination in economic opportunity advertising. *arXiv preprint arXiv:2008.09656* (2020).
- [39] Lauren Kirchner and Matthew Goldstein. 2020. Access Denied: Faulty Automated Background Checks Freeze Out Renters. (2020).
- [40] Ava Kofman and Ariana Tobin. 2019. Facebook Ads Can Still Discriminate Against Women and Older Workers, Despite a Civil Rights Settlement. *ProPublica*. *Retrieved December 12 (2019)*, 2020.
- [41] Ava Kofman and Ariana Tobin. 2021. Facebook Ads Can Still Discriminate Against Women and Older Workers, Despite a Civil Rights Settlement. (2021).
- [42] Sakis Kotsantonis and George Serafeim. 2019. Four things no one will tell you about ESG data. *Journal of Applied Corporate Finance* 31, 2 (2019), 50–58.
- [43] Deepak Kumar, Patrick Gage Kelley, Sunny Consolvo, Joshua Mason, Elie Bursztein, Zakin Durumeric, Kurt Thomas, and Michael Bailey. 2021. Designing Toxic Content Classification for a Diversity of Perspectives. *arXiv preprint arXiv:2106.04511* (2021).
- [44] Issie Lapowsky. 2021. Why Facebook's data-sharing project ballooned into a 2-year debacle. (2021).
- [45] Mark Latonero and Aaina Agarwal. 2021. Human Rights Impact Assessments for AI: Learning from Facebook's Failure in Myanmar. (2021), 18.
- [46] Alexandra Levine. 2021. 'Chilling': Facial recognition firm Clearview AI hits watchdog groups with subpoenas.
- [47] Ada Lovelace and UK DataKind. 2020. *Examining the black box: Tools for assessing algorithmic systems*. Technical Report. Technical report, AdaLovelace Institute, <https://ico.org.uk/media/about/...>
- [48] Mark MacCarthy. 2019. An Examination of the Algorithmic Accountability Act of 2019. *Available at SSRN 3615731* (2019).
- [49] Markey and Matsui. 2021. Senator Markey, Rep. Matsui Introduce Legislation to Combat Harmful Algorithms and Create New Online Transparency Regime.
- [50] Roy Maurer. 2021. New York City to Require Bias Audits of AI-Type HR Technology.
- [51] KATHLEEN MCGRORY and NEIL Bedi. 2020. Targeted. *The Tampa Bay Times* 3 (2020).
- [52] JB Merrill. 2020. Does facebook still sell discriminatory ads. *The Markup* (2020).
- [53] Jacob Metcalf, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, and Madeleine Clare Elish. 2021. Algorithmic impact assessments and accountability: The co-construction of impacts. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 735–746.
- [54] Brent Mittelstadt. 2019. Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence* 1, 11 (2019), 501–507.
- [55] Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, Madeleine Clare Elish, and Jacob Metcalf. 2021. Assembling Accountability: Algorithmic Impact Assessment for the Public Interest. *Available at SSRN 3877437* (2021).
- [56] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [57] Will Oremus. 2021. Facebook keeps researching its own harms – and burying the findings. (2021).
- [58] Barbara Ortutay. 2021. Facebook shuts out NYU academics' research on political ads.
- [59] Michael O'Leary and Warren Valdmanis. 2021. An ESG Reckoning Is Coming. *Harvard Business Review* (2021).
- [60] Nathaniel Persily. 2021. Facebook Hides Data Showing It Harms Users. Outside Scholars Need Access. (2021). Last accessed 25 October 2021.
- [61] Elizabeth Pollman. 2019. Corporate Social Responsibility, ESG, and Compliance. *Forthcoming, Cambridge Handbook of Compliance* (D. Daniel Sokol & Benjamin van

- Root *eds.*, *Loyola Law School, Los Angeles Legal Studies Research Paper* 2019-35 (2019).
- [62] Michael Power. 1997. *The audit society: Rituals of verification*. OUP Oxford.
- [63] Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. 2020. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 469–481.
- [64] Inioluwa Deborah Raji and Joy Buolamwini. 2019. Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial ai products. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 429–435.
- [65] Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, Joon-seok Lee, and Emily Denton. 2020. Saving face: Investigating the ethical concerns of facial recognition auditing. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 145–151.
- [66] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 33–44.
- [67] Dillon Reisman, Jason Schultz, Kate Crawford, and Meredith Whittaker. 2018. Algorithmic impact assessments: A practical framework for public agency accountability. *AI Now Institute* (2018), 1–22.
- [68] Rashida Richardson. 2021. Defining and Demystifying Automated Decision Systems. *Maryland Law Review*, Forthcoming (2021).
- [69] Rashida Richardson, A Cahn, A Kak, A Diaz, A Samant, et al. 2019. Confronting black boxes: A shadow report of the New York City automated decision system task force. *AI Now Institute* (2019).
- [70] Eva Rosen, Philip ME Garboden, and Jennifer E Cossyleon. 2021. Racial Discrimination in Housing: How Landlords Use Algorithms and Home Visits to Screen Tenants. *American Sociological Review* (2021), 00031224211029618.
- [71] Samantha Ross and Gordon Seymour. 2019. The Developing World of Assurance on Sustainability Reporting. unpublished manuscript.
- [72] Piotr Sapiezynski, Avijit Ghosh, Levi Kaplan, Alan Mislove, and Aaron Rieke. 2019. Algorithms that "Don't See Color": Comparing Biases in Lookalike and Special Ad Audiences. *arXiv preprint arXiv:1912.07579* (2019).
- [73] Andrew D Selbst. 2021. An Institutional View Of Algorithmic Impact Assessments. (2021).
- [74] Nathan Sheard. 2021. Banning Government Use of Face Recognition Technology: 2020 Year in Review.
- [75] G Sherwin and E Bhandari. 2019. Facebook settles civil rights cases by making sweeping changes to its online ad platform.
- [76] Mona Sloane. 2021. The Algorithmic Auditing Trap.
- [77] Mona Sloane, Emanuel Moss, and Rumman Chowdhury. 2021. A Silicon Valley Love Triangle: Hiring Algorithms, Pseudo-Science, and the Quest for Auditability. *arXiv preprint arXiv:2106.12403* (2021).
- [78] Jacob Snow. 2018. Amazon's Face Recognition Falsely Matched 28 Members of Congress With Mugshots. <https://www.aclu.org/blog/privacy-technology/surveillance-technologies/amazons-face-recognition-falsely-matched-28> Last accessed 25 October 2021.
- [79] Chandler Nicholle Spinks. 2019. Contemporary housing discrimination: Facebook, targeted advertising, and the Fair Housing Act. *Hous. L. Rev.* 57 (2019), 925.
- [80] Latanya Sweeney. 2013. Discrimination in online ad delivery. *Commun. ACM* 56, 5 (2013), 44–54.
- [81] Nicholas P Tatonetti, Guy Haskin Fernald, and Russ B Altman. 2012. A novel signal detection algorithm for identifying hidden drug-drug interactions in adverse event reports. *Journal of the American Medical Informatics Association* 19, 1 (2012), 79–85.
- [82] Ariana Tobin. 2019. HUD sues Facebook over housing discrimination and says the company's algorithms have made the problem worse. *ProPublica* (March 28, 2019). Available at <https://www.propublica.org/article/hud-sues-facebook-housing-discrimination-advertising-algorithms> (last accessed April 29, 2019) (2019).
- [83] Daniel Trielli, Jennifer Stark, and Nicholas Diakopoulos. 2017. Algorithm tips: A resource for algorithmic accountability in Government. In *Proc. Computation+ Journalism Symposium*.
- [84] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–14.
- [85] James Vincent. 2021. Google is using AI to help users explore the topics they're searching for — here's how. <https://www.theverge.com/2021/9/29/22696268/google-search-on-updates-ai-mum-explained> Last accessed 2 November 2021.
- [86] Paul Voigt and Axel Von dem Busche. 2017. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed.*, Cham: Springer International Publishing 10 (2017), 3152676.
- [87] Kyle Wiggers. 2021. Google trained a trillion-parameter AI language model. <https://venturebeat.com/2021/01/12/google-trained-a-trillion-parameter-ai-language-model/> Last accessed 2 November 2021.
- [88] Christo Wilson, Avijit Ghosh, Shan Jiang, Alan Mislove, Lewis Baker, Janelle Szary, Kelly Trindel, and Frida Polli. 2021. Building and auditing fair algorithms: A case study in candidate screening. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 666–677.
- [89] Eric Wu, Kevin Wu, Roxana Daneshjou, David Ouyang, Daniel E Ho, and James Zou. 2021. How medical AI devices are evaluated: limitations and recommendations from an analysis of FDA approvals. *Nature Medicine* 27, 4 (2021), 582–584.

## A APPENDIX

### REFERENCES FOR EVIDENCE BASE

- [Fig1] 1995. 17 CFR § 201.102. *Code of Federal Regulations: Appearance and practice before the Commission* (1995).
- [Fig2] Bibek Adhikari, James Alm, and Timothy F Harris. 2021. Small business tax compliance under third-party reporting. *Journal of Public Economics* 203 (2021), 104514.
- [Fig3] Federal Aviation Administration. 2022. Aviation Safety Reporting System (ASRS).
- [Fig4] Kristen M Altenburger and Daniel E Ho. 2019. When algorithms import private bias into public enforcement: the promise and limitations of statistical debiasing solutions. *Journal of Institutional and Theoretical Economics* 175, 1 (2019), 98–122.
- [Fig5] Daniel Aobdia, Preeti Choudhary, and Noah Newberger. 2021. The economics of audit production: What matters for audit quality? An empirical analysis of the role of mid-level managers within the audit firm. *An Empirical Analysis of the Role of Mid-level Managers within the Audit Firm (January 28, 2021)* (2021).
- [Fig6] Ian Ayres, Jeff Lingwall, and Sonia Steinway. 2013. Skeletons in the database: An early analysis of the CFPB's consumer complaints. *Fordham J. Corp. & Fin.* L. 19 (2013), 343.
- [Fig7] John M Barrios. 2021. Occupational Licensing and Accountant Quality: Evidence from the 150-Hour Rule. *Journal of Accounting Research* (2021).
- [Fig8] Miroslava Bavorova, Anica Veronika Fietz, and Norbert Hirschauer. 2017. Does disclosure of food inspections affect business compliance? The case of Berlin, Germany. *British Food Journal* (2017).
- [Fig9] Anna Bedford, Martin Bugeja, Samir Ghannam, and Nelson Ma. 2021. The quality of other assurance services supplied by accounting firms: Evidence from independent expert reports. *International Journal of Auditing* 25, 1 (2021), 40–58.
- [Fig10] Victor Manuel Bennett, Lamar Pierce, Jason A Snyder, and Michael W Toffel. 2013. Customer-driven misconduct: How competition corrupts business practices. *Management Science* 59, 8 (2013), 1725–1742.
- [Fig11] Joseph V Carcello, Ann Vanstraelen, and Michael Willenborg. 2009. Rules rather than discretion in audit standards: Going-concern opinions in Belgium. *The Accounting Review* 84, 5 (2009), 1395–1428.
- [Fig12] Stefano Cascino, Ane Tamayo, and Felix Vetter. 2021. Labor market effects of spatial licensing requirements: Evidence from CPA mobility. *Journal of Accounting Research* 59, 1 (2021), 111–161.
- [Fig13] Monika Causholi, Dennis J Chambers, and Jeff L Payne. 2014. Future nonaudit service fees and audit quality. *Contemporary accounting research* 31, 3 (2014), 681–712.
- [Fig14] Heesun Chung, Catherine Heyjung Sonu, Yoonseok Zang, and Jong-Hag Choi. 2019. Opinion shopping to avoid a going concern audit opinion and subsequent audit quality. *Auditing: A Journal of Practice & Theory* 38, 2 (2019), 101–123.
- [Fig15] Bryan K Church, J Gregory Jenkins, Susan A McCracken, Pamela B Roush, and Jonathan D Stanley. 2015. Auditor independence in fact: Research, regulatory, and practice implications drawn from experimental and archival research. *Accounting Horizons* 29, 1 (2015), 217–238.
- [Fig16] Gary Colbert and Dennis Murray. 1999. State accountancy regulations, audit firm size, and auditor quality: An empirical investigation. *Journal of Regulatory Economics* 16, 3 (1999), 267–286.
- [Fig17] Mark L DeFond and Clive S Lennox. 2011. The effect of SOX on small auditor exits and audit quality. *Journal of Accounting and Economics* 52, 1 (2011), 21–40.
- [Fig18] Esther Duflo, Michael Greenstone, Rohini Pande, and Nicholas Ryan. 2013. Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *The Quarterly Journal of Economics* 128, 4 (2013), 1499–1545.
- [Fig19] Robert G Eccles, Michael P Krzus, Jean Rogers, and George Serafeim. 2012. The need for sector-specific materiality and sustainability reporting standards. *Journal of applied corporate finance* 24, 2 (2012), 65–71.
- [Fig20] Matthias Ehng and Harald Hau. 2015. Structured debt ratings: Evidence on conflicts of interest. *Journal of Financial Economics* 116, 1 (2015), 46–60.
- [Fig21] Alex C Engler. 2021. Independent auditors are struggling to hold AI companies accountable. *FastCompany*.

- [Fig22] Corinna Ewelt-Knauer, Anna Gold, and Christiane Pott. 2012. What do we know about mandatory audit firm rotation. *Institute of Chartered Accountants in Scotland (ICAS), Edinburgh. Available at: icas.org.uk/MAFR.pdf* (accessed 22 January 2013) (2012).
- [Fig23] Lily Fang and Ayako Yasuda. 2009. The effectiveness of reputation as a disciplinary mechanism in sell-side research. *The Review of Financial Studies* 22, 9 (2009), 3735–3777.
- [Fig24] U.S. Food and Drug Administration. 2022. Manufacturer and User Facility Device Experience. (2022).
- [Fig25] Pingyang Gao and Gaoqing Zhang. 2019. Auditing standards, professional judgment, and audit quality. *The Accounting Review* 94, 6 (2019), 201–225.
- [Fig26] David A Geier and Mark R Geier. 2004. A review of the vaccine adverse event reporting system database. *Expert Opinion on Pharmacotherapy* 5, 3 (2004), 691–698.
- [Fig27] Brandon Gipper, Luzi Hail, and Christian Leuz. 2021. On the economics of mandatory audit partner rotation and tenure: Evidence from PCAOB data. *The Accounting Review* 96, 2 (2021), 303–331.
- [Fig28] Dan D Goldhaber and Dominic J Brewer. 2000. Does teacher certification matter? High school teacher certification status and student achievement. *Educational evaluation and policy analysis* 22, 2 (2000), 129–145.
- [Fig29] John M Griffin and Dragon Yongjun Tang. 2011. Did credit rating agencies make unbiased assumptions on CDOs? *American Economic Review* 101, 3 (2011), 125–30.
- [Fig30] Elina Haapamäki and Jukka Sihvonen. 2019. Research on International Standards on Auditing: Literature synthesis and opportunities for future research. *Journal of International Accounting, Auditing and Taxation* 35 (2019), 37–56.
- [Fig31] Cassandra Handan-Nader, Daniel E Ho, and Becky Elias. 2020. Feasible policy evaluation by design: A randomized synthetic stepped-wedge trial of mandated disclosure in King County. *Evaluation Review* 44, 1 (2020), 3–50.
- [Fig32] Michael Harber, Ben Marx, and Phillip De Jager. 2020. The perceived financial effects of mandatory audit firm rotation. *Journal of International Financial Management & Accounting* 31, 2 (2020), 215–234.
- [Fig33] Paul M Healy and Krishna G Palepu. 2003. The fall of Enron. *Journal of economic perspectives* 17, 2 (2003), 3–26.
- [Fig34] David Hemenway and Sara J Solnick. 1990. "You Better Shop Around": The Market for Motor Vehicle Inspection. *Law & Policy* 12, 4 (1990), 317–329.
- [Fig35] Kelly J Henning. 2004. What is syndromic surveillance? *Morbidity and mortality weekly report* (2004), 7–11.
- [Fig36] Daniel E Ho. 2012. Fudging the nudge: Information disclosure and restaurant grading. *Yale LJ* 122 (2012), 574.
- [Fig37] Daniel E Ho. 2017. Does peer review work: An experiment of experimentalism. *Stan. L. Rev.* 69 (2017), 1.
- [Fig38] Keith A Houghton and Christine A Jubb. 2003. The market for financial report audits: Regulation of and competition for auditor independence. *Law & Policy* 25, 3 (2003), 299–321.
- [Fig39] Thomas N Hubbard. 2002. How do consumers motivate experts? Reputational incentives in an auto repair market. *The Journal of Law and Economics* 45, 2 (2002), 437–468.
- [Fig40] John Xuefeng Jiang, Mary Harris Stanford, and Yuan Xie. 2012. Does it matter who pays for bond ratings? Historical evidence. *Journal of Financial Economics* 105, 3 (2012), 607–621.
- [Fig41] Barry W Johnson and Peter J Rose. 2019. Federal Tax Compliance Research: Tax Gap Estimates for Tax Years 2011–2013. *Washington, DC: IRS* (2019).
- [Fig42] Stanislav Karapetrovic and Walter Willborn. 2001. Audit system: concepts and practices. *Total Quality Management* 12, 1 (2001), 13–28.
- [Fig43] Hakwoon Kim, Hyoik Lee, Jong Eun Lee, et al. 2015. Mandatory audit firm rotation and audit quality. *Journal of Applied Business Research (Jabr)* 31, 3 (2015), 1089–1106.
- [Fig44] Morris M Kleiner. 2000. Occupational licensing. *Journal of Economic Perspectives* 14, 4 (2000), 189–202.
- [Fig45] Morris M Kleiner and Robert T Kudrle. 2000. Does regulation affect economic outcomes? The case of dentistry. *The Journal of Law and Economics* 43, 2 (2000), 547–582.
- [Fig46] Marielle CJL Klerks, Corry AJ Ketelaars, and Paul BM Robben. 2013. Unannounced, compared with announced inspections: A systematic review and exploratory study in nursing homes. *Health policy* 111, 3 (2013), 311–319.
- [Fig47] Henrik Jacobsen Kleven, Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez. 2011. Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark. *Econometrica* 79, 3 (2011), 651–692.
- [Fig48] Constantine E Kontokosta and Boyeong Hong. 2021. Bias in smart city governance: How socio-spatial disparities in 311 complaint behavior impact the fairness of data-driven decisions. *Sustainable Cities and Society* 64 (2021), 102503.
- [Fig49] Balázs Kovács, David W Lehman, and Glenn R Carroll. 2020. Grade inflation in restaurant hygiene inspections: Repeated interactions between inspectors and restaurateurs. *Food Policy* 97 (2020), 101960.
- [Fig50] Jagan Krishnan, Dasaratha Rama, and Yinghong Zhang. 2008. Costs to comply with SOX Section 404. *Auditing: A Journal of Practice & Theory* 27, 1 (2008), 169–186.
- [Fig51] Phillip T Lamoreaux. 2016. Does PCAOB inspection access improve audit quality? An examination of foreign firms listed in the United States. *Journal of Accounting and Economics* 61, 2–3 (2016), 313–337.
- [Fig52] Bradley Larsen. 2013. Occupational licensing and quality: Distributional and heterogeneous effects in the teaching profession. *Available at SSRN 2387096* (2013).
- [Fig53] Angela Littwin. 2014. Why process complaints: Then and now. *Temp. L. Rev.* 87 (2014), 895.
- [Fig54] Timothy D Lytton and Lesley K McAllister. 2014. Oversight in private food safety auditing: Addressing auditor conflict of interest. *Wis. L. Rev.* (2014), 289.
- [Fig55] Célestin Matte, Natalia Bielova, and Cristiana Santos. 2020. Do cookie banners respect my choice?: Measuring legal compliance of banners from iab europe's transparency and consent framework. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 791–809.
- [Fig56] Patricia A McCoy. 2002. Realigning Auditors' Incentives. *Conn. L. Rev.* 35 (2002), 989.
- [Fig57] Anthony D Miyazaki. 2008. Online privacy and the disclosure of cookie use: Effects on consumer trust and anticipated patronage. *Journal of Public Policy & Marketing* 27, 1 (2008), 19–33.
- [Fig58] Don A Moore, Lloyd Tanlu, and Max H Bazerman. 2010. Conflict of interest and the intrusion of bias. *Judgment and Decision Making* 5, 1 (2010), 37.
- [Fig59] Kevin J Murphy and Tatiana Sandino. 2010. Executive pay and "independent" compensation consultants. *Journal of Accounting and Economics* 49, 3 (2010), 247–262.
- [Fig60] NIST. 2022. National Vulnerability Database.
- [Fig61] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. 2020. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [Fig62] Dara O'Rourke. 2000. Monitoring the monitors: A critique of PricewaterhouseCoopers (PwC) labour monitoring. *Department of Urban Studies and Planning Massachusetts Institute of Technology* (2000).
- [Fig63] Lamar Pierce and Michael W Toffel. 2013. The role of organizational scope and governance in strengthening private monitoring. *Organization Science* 24, 5 (2013), 1558–1584.
- [Fig64] Scott Pluta. 2013. So, how do I submit a complaint?
- [Fig65] Katherine Porter. 2012. The Complaint Conundrum: Thoughts on the CFPB's Complaint Mechanism. *Brook. J. Corp. Fin. & Com. L.* 7 (2012), 57.
- [Fig66] Toshiyuki Sakaeda, Akiko Tamon, Kaori Kadoyama, and Yasushi Okuno. 2013. Data mining of the public version of the FDA Adverse Event Reporting System. *International journal of medical sciences* 10, 7 (2013), 796.
- [Fig67] Iskander Sanchez-Rola, Matteo Dell'Amico, Platon Kotzias, Davide Balzarotti, Leyla Bilge, Pierre-Antoine Vervier, and Igor Santos. 2019. Can i opt out yet? gdpr and the global illusion of cookie control. In *Proceedings of the 2019 ACM Asia conference on computer and communications security*. 340–351.
- [Fig68] Peter Schuck. 1972. curious case of the indicted meat inspectors. *Harper's magazine* (1972).
- [Fig69] Jodi L Short, Michael W Toffel, and Andrea R Hugill. 2016. Monitoring global supply chains. *Strategic Management Journal* 37, 9 (2016), 1878–1897.
- [Fig70] Thair Htut Soe, Oda Elise Nordberg, Frode Guribye, and Marija Slavkovik. 2020. Circumvention by design-dark patterns in cookie consent for online news outlets. In *Proceedings of the 11th nordic conference on human-computer interaction: Shaping experiences, shaping society*. 1–12.
- [Fig71] David G Stevenson. 2006. Nursing home consumer complaints and quality of care: A national view. *Medical Care Research and Review* 63, 3 (2006), 347–368.
- [Fig72] Lior Jacob Strahilevitz. 2006. How's My Driving-For Everyone (and Everything). *NYUL Rev.* 81 (2006), 1699.
- [Fig73] Nicholas P Tattoneti, Guy Haskin Fernald, and Russ B Altman. 2012. A novel signal detection algorithm for identifying hidden drug-drug interactions in adverse event reports. *Journal of the American Medical Informatics Association* 19, 1 (2012), 79–85.
- [Fig74] William R Titera. 2013. Updating audit standard—Enabling audit data analysis. *Journal of Information Systems* 27, 1 (2013), 325–331.
- [Fig75] Veronica Toffolatti, Martin McKee, and David Stuckler. 2017. Evidence points to 'gaming' at hospitals subject to National Health Service Cleanliness Inspections. *Health Affairs* 36, 2 (2017), 355–361.
- [Fig76] Martino Trevisan, Stefano Traverso, Eleonora Bassi, and Marco Mellia. 2019. 4 Years of EU Cookie Law: Results and Lessons Learned. *Proc. Priv. Enhancing Technol.* 2019, 2 (2019), 126–145.
- [Fig77] Jerry L Turner, Theodore J Mock, Paul J Coram, and Glen L Gray. 2010. Improving transparency and relevance of auditor communications with financial statement users. *Current issues in auditing* 4, 1 (2010), A1–A8.
- [Fig78] Briana Vecchione, Karen Levy, and Solon Barocas. 2021. Algorithmic Auditing and Social Justice: Lessons from the History of Audit Studies. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–9.

Audit Ecosystem Design Feature	Explanation
Mandatory?	Is the audit required by law or voluntarily undertaken?
Public or private?	Is the audit conducted by the public sector (e.g. a government agency) or by a private entity?
Accreditation requirements?	Does the auditing entity have to be accredited, certified, or recognized by an official body before it is allowed to conduct the audit?
What is being audited?	What is the subject of the audit? Is it a product? Compliance with certain regulations?
Can the auditor make recs/changes?	Does the auditor have the power to make recommendations or changes to the audited entity?
Are audits recurring?	Does the audit take place just once? Or is there a periodic auditing requirement?
Who selects the auditor?	Who selects the entity that will conduct the audit?
Who pays the auditor?	Who pays the entity that will conduct the audit?
Public transparency requirements?	Is there a requirement that the contents of the audit be made available to the public?
Rotation requirements?	Is there a requirement that the auditors rotate off of the auditing process after a certain period of time?

Table 1: Analyzed Audit Ecosystem Features

Auditing program	Description
<b>First Party Audits</b>	
EPA Audit Policy	The Environmental Protection Agency (EPA) issued its "Audit Policy" in 1995. The Policy, which is designed to increase compliance with environmental regulations, eliminates or reduces civil penalties for violations that facilities discover as a result of a self-audit and voluntarily disclose.
SQA Audits	Software Quality Assurance (SQA) audits are conducted throughout the life cycle of a software program to quality-check the engineering processes and methods used in the software's development. SQA audits are typically conducted in-house, and can include various audit types, including security audits and information systems audits.
<b>Second Party Audits</b>	
Supply chain audits	Because the specific features of second-party audits vary depending on the contractual terms between the auditing and audited entities, this is a case study of a two-year audit of suppliers in the confectionary industry. See Djekic et al. [17].
<b>Third Party Audits</b>	
SOX financial audits	The Sarbanes-Oxley Act (SOX) requires that companies hire independent auditors to complete yearly checks of their financial reports. A SOX compliance audit also reviews companies' internal controls over financial reporting (e.g. management systems; data security practices) to make sure they are secured against fraud or theft.
MDSAP	The Medical Device Single Audit Program (MDSAP) allows medical device manufacturers to satisfy the requirements of multiple regulatory jurisdictions with a single regulatory audit of the medical device. The MDSAP certification process takes three years, during which participants undergo annual audits.
FDA audits of clinical trial data	The U.S. Food and Drug Administration (FDA) routinely audits clinical trials to confirm the validity of their findings, including trials that may form the basis for FDA drug approval. The FDA also conducts for-cause audits of investigators if someone reports that investigator to the FDA or if the FDA receives a complaint about the investigator from the drug manufacturer.
California GHG Verification	According to California state regulations, greenhouse gas (GHG) emissions data reports must be independently checked every year by verification bodies accredited by the California Air Resources Board (CARB).
USDA National Organic Program	Farms and businesses that want to label a product as "organic" must be certified by an accredited third-party organization. The audit is meant to ensure that the product meets national organic standards.
CPSC third-party testing program for children's products	Under federal law, every toy intended for use by children age 12 or under must be tested by an accredited third-party laboratory for compliance with the applicable children's product safety requirements.
FDA FAERS post-market surveillance	Under U.S. Federal Communication Commission (FCC) regulations, Telecommunication Certification Bodies (TCB) must audit samples of telecommunications products after they have gone to market (i.e. as part of a post-market surveillance program) to ensure that the products remain in compliance with FCC requirements.
FCC TCB post-market surveillance	The FDA Adverse Event Reporting System (FAERS) is a computerized database that assists the FDA with its postmarket surveillance activities. Patients, caregivers, and healthcare professionals may submit adverse drug events to FAERS. The reports are assessed by a multidisciplinary staff (and in some cases, by AI/ML tools), after which the FDA may conduct further postmarket studies or take regulatory actions (e.g. updating labels or even reevaluating an approval decision) to improve product safety.
FAA Office of Audit and Evaluation	The U.S. Federal Aviation Administration (FAA)'s Office of Audit and Evaluation (AAE) receives protected disclosures via its online hotline from FAA employees, certificate holders, and employees of certificate holders concerning potential violations of the FAA's statute and regulations and acts, omissions, or gross misconduct that affect aviation safety. AAE evaluates these disclosures, and if it determines that further investigation is warranted, it will either refer the disclosure to another FAA office or request assistance to conduct an inquiry.
HHS nursing home inspections	To participate in the Medicaid and Medicare programs, nursing homes must undergo annual health inspections. These inspections are conducted by state survey agencies, and inspections may take place more often if the nursing home is performing poorly, or if there are complaints or reported incidents within the facility.

Table 2: Overview of Select U.S. Audit Programs