

Auditing Work: Exploring the New York City algorithmic bias audit regime

Lara Groves
lgroves@adalovelaceinstitute.org
Ada Lovelace Institute
United Kingdom

Jacob Metcalf
jake.metcalf@datasociety.net
Data & Society Research Institute
USA

Alayna Kennedy
alayna.a.kennedy@gmail.com
Independent Researcher
USA

Briana Vecchione
briana@datasociety.net
Data & Society Research Institute
USA

Andrew Strait
AStrait@adalovelaceinstitute.org
Ada Lovelace Institute
United Kingdom

ABSTRACT

In July 2023, New York City (NYC) implemented the first attempt to create an algorithm auditing regime for commercial machine-learning systems. Local Law 144 (LL 144), requires NYC-based employers using automated employment decision-making tools (AEDTs) in hiring to be subject to annual bias audits by an independent auditor. In this paper, we analyse what lessons can be learned from LL 144 for other national attempts to create algorithm auditing regimes. Using qualitative interviews with 17 experts and practitioners working within the regime, we find LL 144 has failed to create an effective auditing regime: the law fails to clearly define key aspects like AEDTs and what constitutes an independent auditor, leaving auditors, vendors who create AEDTs, and companies using AEDTs to define the law's practical implementation in ways that failed to protect job applicants. Several factors contribute to this: first, the law was premised on a faulty transparency-driven theory of change that fails to stop biased AEDTs from being used by employers. Second, industry lobbying led to the definition of what constitutes an AEDT being narrowed to the point where most companies considered their tools exempt. Third, we find auditors face enormous practical and cultural challenges gaining access to data from employers and vendors building these tools. Fourth, we find wide disagreement over what constitutes a legitimate auditor and identify four different kinds of 'auditor roles' that serve different functions and offer different kinds of services. We conclude with four recommendations for policymakers seeking to create similar bias auditing regimes that use clearer definitions and metrics and more accountability. By exploring LL 144 through the lens of auditors, our paper advances the evidence base around audit as an accountability mechanism, and can provide guidance for policy-makers seeking to create similar regimes.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

FAccT '24, June 03–06, 2024, Rio de Janeiro, Brazil
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0450-5/24/06
<https://doi.org/10.1145/3630106.3658959>

CCS CONCEPTS

• **Computing/technology policy**; • **Social and professional topics** → **Governmental regulations**; **Government technology policy**;

KEYWORDS

algorithm audit, algorithmic bias, AI policy, Local Law 144

ACM Reference Format:

Lara Groves, Jacob Metcalf, Alayna Kennedy, Briana Vecchione, and Andrew Strait. 2024. Auditing Work: Exploring the New York City algorithmic bias audit regime. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24), June 03–06, 2024, Rio de Janeiro, Brazil*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3630106.3658959>

1 INTRODUCTION

In July 2023, New York City (NYC) became the first jurisdiction to implement a law that mandates independent algorithmic bias audits for commercial companies and city agencies, specifically focusing on automated employment decision-making tools (AEDTs) in hiring and promotion. Under this law—known as Local Law 144 (LL 144)—all NYC-based employers using AEDTs are obligated to hire a third-party independent auditor to conduct annual 'bias audits' and post the resultant audit reports on their website. Additionally, the law requires that those employers provide jobseekers with a transparency notice regarding their use of AEDTs and the right to opt-out of analysis by the algorithmic system in favour of a human decision process when otherwise required by law (e.g., disability accommodations).

LL 144 is the first law to create a third-party algorithm audit regime for AI and machine-learning systems. Algorithm audits are an emerging method for assessing an AI system for a particular kind of legal or ethical risk. They aim to create greater accountability for AI system developers [26] or inscribe normative standards[33]. Audits can encompass a range of different practices and assess for different risks [9]. This includes technical audits of a system's inputs and outputs to determine if the system performs differently for different user groups [47], auditing whether a system complies with local regulation or internal standards of development, and sociotechnical audits that assess how a system is impacting wider societal processes and contexts [45]. Auditing methodologies stem primarily from the fields of computer science

and data science, and have largely focused on audits assessing issues of bias in AI systems[10]. Previous research has highlighted that audits conducted by an independent third party tend to be of higher quality than audits conducted by an internal team or a contracted second-party assessor [7, 23, 54]. In other sectors, like finance or environmental studies, third-party auditing regimes require underlying governance mechanisms—such as transparency requirements to publish the audit, standards of auditor practice, and an oversight body to adjudicate instances of malpractice. These create a functional auditing ecosystem in which multiple actors can evaluate, test and audit AI systems before, during and after they are deployed[48]. Existing theory and practice has informed a growing literature of algorithm audits, which also identify the need for clearly defined standards, mandatory disclosure of reports, and formalisation of algorithm auditor practice, including access powers and accreditation[5, 8, 21, 48, 56].

While LL 144 is the first law to create an algorithm auditing regime, there is some history of regulators and government agencies using algorithm audits to assess bias in AI systems. Some national regulators in regions like Australia ACCC [19] and the Netherlands [49] have the power to audit algorithms to assess their legality or compliance with national regulation or law, including issues of algorithmic bias. An increasing number of global AI governance policy proposals are now seeking to establish auditing regimes as part of a wider AI governance process. The European Union’s Digital Services Act EC [15] which comes into force in February 2024, requires ‘very large online platforms’ to conduct regular audits of their compliance with the law, which may often include audit practices similar to LL 144. Similar auditing requirements are also being proposed in legislation in the United Kingdom [42] and the United States Congress [50]. The US state of Colorado also passed a statute to prevent bias in the insurance industry through commercial algorithmic systems shortly after LL 144, which went into effect November 2023, although that does not require independent audits[38, 43]. The implicit theory of change animating LL 144 is an assumption common in policymaking: by incentivising and supporting a market for assessment and transparency documentation, assessors will be able to propagate best practices and protect civil rights even without strict regulatory limits on the design of technical systems. Additionally, by providing the public—in this case primarily jobseekers, and secondarily civil society and academic researchers—with insight into the algorithmic systems making impactful decisions about their lives, the public will make better informed decisions and push the vendors building these tools and the companies using them to create fair(er) systems. In this study, we explore the experience of those conducting independent audits to illustrate whether the particular accountability structures imposed by this first-of-its-kind law will achieve those goals. We centre the experience of auditors in this study because they are the node who must interpret the needs and interests of employers, developers, regulators, and the jobseekers the law is meant to protect, all while providing a financially viable service. Using qualitative interviews with 17 experts and practitioners working within the LL 144 bias audit regime, our paper seeks to answer three research questions (RQs):

- **RQ1:** What are the practical components of a bias audit in this context?
- **RQ2:** What are the relational dynamics and incentives that make for an effective bias auditing regime?
- **RQ3:** What are the experiences of auditors, and how can they inform wider policy and practice?

The core finding of our study is that the work of algorithm auditing in this regime is largely about managing the relational dynamics between stakeholders established by the accountability structures in the law. This paper offers novel empirical evidence into an emerging bias audit regime with the aim of informing wider policy and practice around algorithm audits and auditors at the local, national, and international level. To our knowledge, this is the sole study exploring the LL 144 regime through the lens of auditors and one of the first attempts to use a case study of a bias auditing regime to discern lessons for other attempts. By developing the evidence base around audits, we can draw stronger conclusions about auditing as a potential accountability mechanism and contribute to an ecosystem of safe and ethical AI systems in high-stakes decision-making domains like recruitment.

2 METHODS

2.1 Expert interviews

Recruitment for interviews followed a combination of direct and snowball methods: we identified auditors via the publicly available audits, through contacting auditors in our networks, and from recommendations from our interviewees. We sent a total of 42 invitations. The interviews were conducted primarily by the lead author, with support and contributions from other authors, and took place from July 2023 to January 2024. All interviews took place virtually, using video conferencing software, and were transcribed using a speech-to-text transcription service. One interviewee did not consent to their interview quotes being used in this paper.

2.2 Data analysis

We adopted a grounded theory epistemology, allowing us to surface latent themes, patterns, and social interactions from our data [13]. Four authors contributed to the coding process using an inductive approach, allowing us to derive codes based on concepts from the data. We used Atlas.ti data analysis software to support the process. We completed a round of asynchronous coding, generating an initial codebook of 102 codes. We then deliberated and identified areas of consensus, dissensus and overlap, enabling a synthesis of the codebook down to 16 codes, represented in Table 1. Each interview was then coded again by two separate authors using this revised codebook.

Auditing service under LL144	Auditor legitimacy
Background of auditor	Beneficiaries of the LL 144 regime
Comparison with other (audit) domains	Components of the audit
Data practices/issues/metrics	Enforcement of LL144
Evaluating audit/auditor practice	History/context of the law
Incentives to comply with the law	Lessons for policymakers
Opinion on compliance	Opinions of the law

Table 1: Codes

2.3 Limitation of study

This paper reports on the implementation of LL 144 in NYC, and its relationship to United States federal employment law and civil rights law. This gives our research a narrow geographical focus. We believe that there are lessons emerging from our findings that will nonetheless be of utility to other jurisdictions/domains. A limitation for our interview findings is that not every auditor with attributable published audits under LL 144 accepted our request to participate. Additionally, the management consulting firms currently offering AI audits for compliance with other legislation targeting AI/data-driven systems (e.g. the European Union’s Digital Services Act) are not conducting audits under the LL 144 regime at the time of this research and therefore are not represented in this study. Participants were not offered compensation for their participation.

3 BACKGROUND OF LOCAL LAW 144

An introduction to the specifics of LL 144 is critical to engaging with our thesis. In this section, we present core definitions, the central tenets of the law, and insight into the history and context in which the law emerges. Significantly, LL 144 offers particular definitions of key terms; the same terms may be defined differently in other contexts.

3.1 History of the law

LL 144 is the result of a long-running effort by the municipal government of NYC, pushed by many civil rights advocates, to create a regulatory structure for algorithmic systems [11, 31, 34]. This effort first focused on city government systems and eventually regulating only commercial algorithmic hiring systems. It was first passed by the NYC Council in late 2021 [22] which tasked the Department of Consumer and Worker Protection (DCWP) with rule-making and enforcement. There were two revisions, with a period of open public comment, and multiple delays to the lawmaking process. The third and final version was adopted in April 2023 and implemented in July 2023[37]. See Figure 1 ‘History of the law’ for more details.

3.2 Key aspects of the law and its context in US anti-discrimination law

While LL 144 marks the first attempt to require bias auditing of AI systems, its text and theory of change awkwardly interacts with existing US employment anti-discrimination law. Some of the key aspects of this law are borrowed from existing litigation and standards—for example, definitions for race and gender and established formulas for assessing bias through impact ratios, as we detail below. Some, such as the definition of AEDT, relied on law drafts, public comments and revisions to shape. We found that auditors were often involved in negotiating how exactly the novel definitions offered by LL 144 should be interpreted in light of established norms and laws.

LL 144 requires only employers/hiring agencies that use the system to conduct a bias audit, but not the developer/vendor/platform that builds or sells the system [36]¹. The law imposes two primary

¹The DCWP’s FAQ states that “Employers and employment agencies are responsible for ensuring they do not use an AEDT unless a bias audit was done. The vendor that created the AEDT is not responsible for a bias audit of the tool” [36]. The decision to only impose obligations on end-users was primarily due to the city’s jurisdictional

obligations on any NYC employer using AEDTs for decisions on hiring and/or promotion.

- **Annual ‘bias audit’ of the AEDT conducted by an independent auditor** for race and gender features and the intersection thereof. The bias audit report must then be posted on the employer’s website. The prescribed audit is more accurately described as a ‘disparate impact’ audit [5, 6, 24], as it measures a specific form of algorithmic bias [36]. The law defines independent auditors as third-party experts who have no financial stake in the success of the product or the financial outcome of the employer.
- **Transparency notice about use of AEDT.** For job applicants, this disclosure is typically appended to the job listing on the employer’s website and/or recruitment platform(s) along with other disclosures, but can be delivered directly to the jobseeker. LL 144 does not mandate the right to opt-in to an alternative human-only decisions process for everyone, but it does require that if such an alternative is offered to fulfil other civil rights protections (e.g., disability accommodations) then the notice must instruct jobseekers on how to request it.

LL 144 defines *race and gender* using categories established by the EEOC [17]. Other protected categories central to AI fairness scholarship and employment civil rights law are not included in the audit requirement, such as disability, age and religious status. The law requires the audit report to present an impact ratio table. Figure 3 in the Appendix provides an example of these impact ratio tables. The DCWP is tasked with enforcing these rules and assessing civil penalties between \$500 and \$1,500 per day for violations². The jurisdiction of the law covers any position with a primary workplace in the boroughs of NYC, including remote work jobs that are headquartered there.

The law defines AEDTs so that many (perhaps most) of systems one might identify as ‘automated employment decision tools’ on the market and offered by recruitment platforms are out of scope; it also grants employers significant discretion to decide if their systems are in scope. This creates a tension that auditors must help their clients navigate: is *this* AEDT the *right kind* to fall within scope? AEDTs are defined as machine-learning/AI systems that “substantially assist or replace discretionary decision making” in the hiring or promotion process [32, 33]. *Substantially assist* is defined as being the primary/majority reason for a ‘selection’ decision. To the authors’ best knowledge, there is no system on the market that fully replaces human decision-making in hiring, and therefore the scope of the law is *de facto* determined by the meaning of ‘substantially assist.’ However, ‘assisting’ is a feature of the human and organisational aspects of the hiring process, and not of the technical system.

Because of this, two employers using the same model from the same vendor may have different interpretations of whether their

limits—developers of AEDTs or hiring platforms are unlikely to be domiciled in New York City. Additionally, employment law generally focuses on outcomes of decision-making, such as moments of hiring and firing, which is typically undertaken by employers and not by the platforms.

²One auditor noted the enforcement penalties have some ambiguity about whether this fine applies per-candidate, which would drastically increase the cost, or per audit [P10].

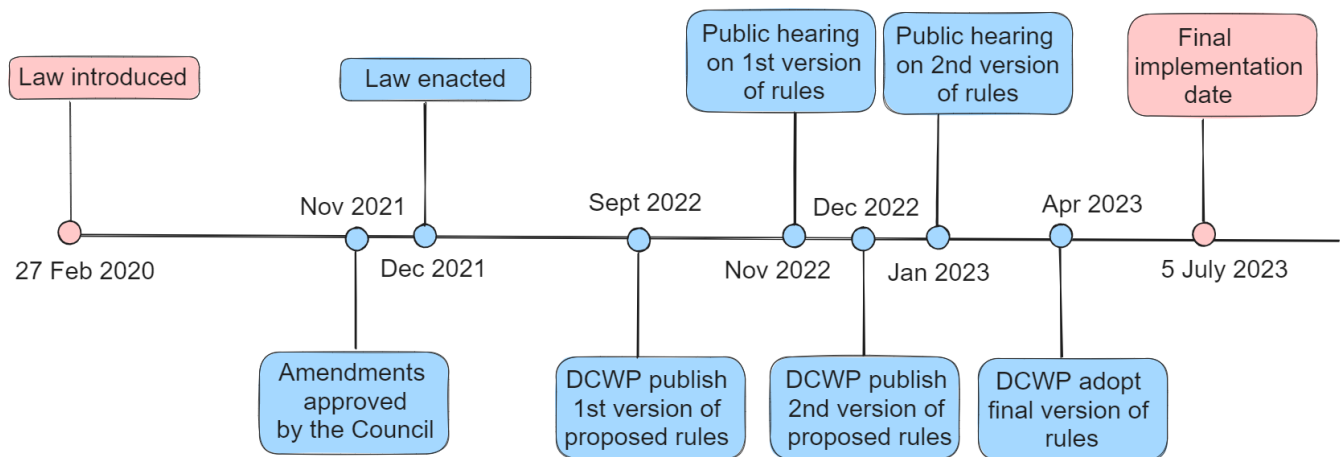


Figure 1: History of LL 144

own output meets LL 144’s definition of ‘substantial assistance’. *Machine learning/artificial intelligence* is defined in terms of outputting a simplified score or ranking on the basis of multiple feature inputs, in which the machine in part algorithmically weights those features. This narrows the scope further to deep-learning methods, and excludes models with even intentionally discriminatory weights that are set manually.

The law draws on concepts from US employment law to determine how an auditor should assess for bias, but also diverges from those laws in ways that create further tensions for auditors to navigate. The law requires the creation of an impact ratio table, where an *impact ratio* is a method for measuring discriminatory outcomes as the relative selection rate between demographic groups. A *selection rate* is the frequency at which members of a group are chosen to move forward in a hiring/promotion process or rejected/screened out; ‘selection’ does not refer only to the final hiring decision, but also all decisions prior. LL 144 also makes use of a ‘scoring rate’ to measure impact ratios, capturing the fact that AEDTs are often used for producing simplified scores or rankings such as personality or intelligence scores. The *scoring rate* considers the frequency at which the mechanism gives members of a group higher than median scores. Therefore, an impact ratio is a comparison of rates (not of absolute numbers), with the rate for one group as the numerator and the rate for the most-selected group or full population as the denominator.³ An impact ratio of 1.0 means a perfectly equal selection rate between groups; an impact ratio less than one indicates a discriminatory outcome against the less-selected group, and the lower the fraction the more discriminatory the outcome is.

Disparate impact (aka adverse impact) is a state of impermissible discriminatory outcomes in US employment law. Disparate impact does not require intent to discriminate; rather, it captures systemic discrimination for which the employer is nonetheless responsible,

³In practice, ‘less-selected’ group typically corresponds to a ‘historically-disadvantaged’ or ‘protected’ category. However, those are not synonymous. The less-selected group could be a historically-advantaged/non-protected group, e.g., if an employer only hires members of a historically-disadvantaged group for low-status jobs. Thus it is important to recognize that disparate impact cannot measure historical discrimination.

measured only by the outcomes of the system[52]⁴. While there are potentially many different types of ‘fairness audits,’ the audits required by LL 144 are most accurately described as ‘disparate impact audits.’ The so-called ‘four-fifths rule’ is a US anti-discrimination convention or norm, where disparate impact is defined as an impact ratio below 0.8 (four-fifths), meaning the selection rate of the less selected-group is below 80% of that of the most-selected group (or entire population) [4]. Rather than a formal regulation, the four-fifths rule operates as a core guideline in the EEOC’s decision tree for permissible selection procedures [16], becoming ingrained in hiring practices over time. An impact ratio that falls below 0.8 is potentially subject to regulatory scrutiny and often must have additional justification to be legal. An impact ratio above 0.8 typically grants a presumption of non-discriminatory outcomes, but is not an absolute shield from discrimination claims—other aspects of the employer’s selection process may nonetheless be illegal. As critics point out, the four-fifths rule is an arbitrary bureaucratic compromise and not grounded in empirical evidence, and so may not be useful for judging algorithmic bias despite its early prominence in the field [18, 55]. Nonetheless, as a prevalent convention, AEDT developers and employers who use AEDTs are strongly incentivised to operate as if it were a rule to minimise legal risk.

Crucially for auditors, LL 144 sets no minimum for permissible impact ratio rates, and no obligations to cease using or repair any system that is found to cause disparate impact—all that is required is publication of the audit [36]. *The primary role the four-fifths rule plays in LL 144 is therefore its absence and the challenging accountability relationships that creates.* Despite the fact that LL 144 makes no cross-reference to the four-fifths rule, in most of our interviews (P4, P5, P6, P7, P8, P9, P10, P13 and P15) auditors cited it as a flawed target that they were obliged to address to satisfy the needs of their clients. Part of LL 144’s implicit theory of change is that creating a market for independent algorithm auditing and transparency reporting will incentivise gradual adoption of better practices, thereby protecting jobseekers’ civil rights. This theory

⁴That no intent need be inferred to measure discriminate outcomes is why disparate impact is often considered illustrative of algorithmic bias [24, 51]

of change relies on transparency, jobseeker autonomy, and reputational pressure, rather than imposing limits on companies or vendors using AEDTs or proactive investigations by regulators (P4, P5). While an algorithm auditor might conceive of their role as helping clients reduce discriminatory outcomes and protect civil rights—via services such as consulting, model design advice, governance practices, legal guidance, etc.—their mandated role is only the production of audit reports. Absent a mandate to remediate discriminatory outcomes, auditors must find secondary avenues to achieve goals beyond transparent reporting and provide higher margin services.

3.3 Analysis of the public comments

As part of our initial background research, we conducted a qualitative analysis of the third and final round of public comments on the DCWP's rule-making process, submitted in January 2023 [3, 7]. Our analysis revealed disagreements primarily about the scope of the law (which systems would be subject to it) and the positionality of auditors (does the legitimacy of audits require the independence of auditors?). Of the 46 total comments, we identified 13 civil society organisations (including some authors of this paper), four audit service firms, six employment law firms, seven system developers, and three employment platforms. Labour rights advocacy organisations, including labour unions, were not well represented, with the bulk of the civil society comments focusing on issues of algorithmic accountability. Only a few comments specifically detailed the consequences for jobseekers of the details in the law.

Generally, both hiring platforms and AEDT vendors expressed concern about the law being too broad and onerous, with civil society stating the reverse, arguing industry actors may try to skirt regulation. We found evidence that AEDT vendors lobbied the government for narrower definitions, language that allowed for exceptions, and higher degrees of control over the audit process. One platform suggested removing the requirement for an independent auditor in favour of an internal team, while a vendor argued the 10-day transparency notice requirement would burden NYC employers. Audit service/audit tooling firms were in favour of independent auditing on grounds of the trustworthiness of audits and market incentives. In official comments and reporting, civil society organisations were split on whether LL 144 should be revised or scrapped. One local coalition opposed the law due to its limitation to hiring decisions, calling for regulatory scrutiny on other forms of algorithmic employment tools, like work assignments or discipline [11]. Civil society commenters frequently called for the removal of the 'substantially assist' criterion, arguing it would allow employers to declare themselves out of scope—effectively nullifying the law. In contrast, industry commenters stated that it should remain, which it did. The tensions between groups in the public comments highlights the role that relational dynamics between industry, civil society, and the DCWP played in shaping the final law.

3.4 Public availability of audit reports

Despite the law's emphasis on transparency, very few audits are available publicly in the first six months of final rule implementation. We searched for available audits using ad hoc search methods prior to conducting interviews, we identified 19 audits claiming to

be related to LL 144, including five from system developers/vendors not subject to audit under the law. Some of these audits are no longer available online. Both follow-up systematic searches four months after implementation in a companion project [57] and an open-source collection assembled by researchers at the New York branch of the American Civil Liberties Union [25] found commensurate results. Surveys indicate that AEDTs of all varieties are thought to be widespread, but there is no available data about the baseline rate of adoption of AEDTs generally, let alone adoption of systems that fall within the scope of LL 144 [12, 28, 52]. Therefore, it is not possible to know precisely how many audits we should expect to find, and lacking a central repository, it is also impossible to know precisely how many audits are published. However, given the many thousands of employers in NYC and widespread use of AEDTs, one would reasonably anticipate that more than a handful of employers would fall under scope of this law.

Under LL 144, employers have almost-total discretion to make a judgement on whether their systems are in scope for multiple reasons, and jobseekers have no formal mechanism to challenge that decision. The absence of an audit on an employer's website neither implies the absence of an AEDT (as they may use an AEDT self-determined as out of scope), nor non-compliance (they may use a genuinely out of scope AEDT or none at all), nor even a lack of audit (the process may have been conducted, but the report is not published). Though the audit reports themselves were not a focal point for our study, mapping the audits and observing these conditions was informative for our interviews with auditors.

4 FINDINGS FROM INTERVIEWS

Our initial investigation into the public comments that shaped the law revealed relational tensions and potential gaps in the definitions created by the law. While public comments were split on whether the law was too narrow, or too broad, we wanted to gather information from auditors about the actual implementation of the law. Informed by our investigation into the law, we asked our interview participants questions roughly following four themes:

- **High-level organisational questions**, including questions on the auditor team and their experience
- **Relational dynamics** between auditor, companies audited, AEDT tool developers and regulators, including questions on how key information is communicated
- **Audit methodologies**, including questions on the content, process, artefact of audit and success criteria
- **Learnings from the audit process or broader LL 144 regime** for use by policymakers

See Table 2 in the Appendices for Participant ID, including background and organisational information, and 4.1 for the kinds of roles and services auditors enact in the LL 144 regime

4.1 Opinions of the law overall were that it was well-intentioned but seriously flawed

'I don't think [LL 144] should be admonished for not being perfect. I call this the first pancake: without the first pancake, none of the other pancakes would be better. And the first one is always pretty awful in the pan, no matter what you do.' [P6]

We asked our interviewees for their opinions of the law and the law-making process, which revealed a consensus about some fundamental challenges that impeded auditing practices. While most auditors we interviewed had a positive opinion about the spirit of the law, they also agreed that it was not a particularly effective auditing mechanism. Participants gave strong opinions on employers' discretion to decide if systems meet the 'substantially assist' definition (see 2.3): '[the AEDT definition] creates huge loopholes'. Participants with knowledge of the law-drafting process pointed to lobbying efforts by both AI vendors and large employers to water down the definition. Many auditors [P2, P5, P14] noted a key challenge with the legislation is that "it's too open for interpretation when it comes to the scope of the systems" [P2]. As [P5] put it, the "first hurdle" many auditors experienced with this law was their clients asking if it was a good use of resources. Some auditors did not like the law's theory of change. They expressed a view that it does not do enough to protect job applicants because audits and notices to applicants, while important, will not enable them to make an informed decision about whether to be subject to an AEDT or a human reviewer, or about whether to apply for the job at all [P11, P13]. Others [P1] noted that the audit reports can serve as the basis for civil lawsuits or regulatory action against a company or vendor if an audit shows the tool may be biased against a certain protected demographic, but P4 noted that this theory of change doesn't account for the impossible burden of proof an injured jobseeker would face in making a successful complaint to the DCWP. In their view, the burden of proof for discrimination needs to be revised for this theory of change to work.

4.2 Auditing regimes create a variety of auditing roles and services

Our interviews reveal little shared agreement of the role and remit of an 'auditor' of AEDT systems in this regime, and as an algorithmic auditor more generally. Rather than identifying one single function or role of an auditor, we found participants broadly self-identified into one of four categories of audit 'actors' that offered different services (or a mixture of these services) to help clients comply with LL 144:

- Participants that offered a 'pre-audit' service to help a company become 'audit ready,' including helping with their data collection process, providing data governance tools, or other frameworks [e.g. P1, P5, P6]
- Participants that conducted the audit and wrote the report for a client [P2, P3, P7, P8, P11, P12]
- Participants that offered additional guidance and mitigation strategies for the company [P9, P10, P13, P14, P15]
- Participants who offer a service to certify that an audit has been conducted in an appropriate manner [P11, P12]

Clients received different services when conducting an audit, depending on the auditor. For example, participants from law firms [P3, P16] might offer a judgement about whether the client is compliant with the law and offer strategies to become compliant. Additionally, while the law plainly imposes obligations on only employers, the DCWP does offer secondary guidance that vendors 'may' conduct an independent audit for their own purposes (presumably for market advantage)—a phenomenon observable in the various

collections of audit reports—which means that auditors must also consider the needs of clients with different regulatory statuses [36]. Differences in services can also be observed in the public audits artefacts, which vary widely between the minimal quantitative measures of impact ratios required by the law, to more extensive assessments of accountability processes and audit readiness.⁵ The lack of clear roles, responsibilities, and services around the law furthers the need for additional relational work to comply with the law. Interpretation and implementation are left to auditors and their clients, resulting in fractured roles and responsibilities.

4.3 Disagreement over what constitutes a legitimate auditor and legitimate auditing practice

'I think it's important that we are not auditors, because auditors by our interpretation of that definition, are an independent authority ... nothing that an auditor does should be taken as advice. We see them as an assessor that can validate.' [P6]

The algorithmic auditing ecosystem is nascent and currently lacks the structure and standards of other audit domains, such as financial services [48]. Auditors who provide services to LL 144 clients are both interpreting the context of that law, and conceptualising what a sustainable algorithmic auditing business may look like in the long-run—how might they provide multiple forms of utility to their clients and persuade potential future clients of their value? This relational dynamic often strains the traditional meanings of 'audit.' Participants expressed different views on what practices or qualities made them legitimate auditors under LL 144. One interviewee cited that the diverse input involved in their "crowdsourced, participatory" [P1] audit certification service would raise legitimacy questions, while another stated their "transparent approach" [P12] defined them as a legitimate auditor practice.

Interviewees cited their domain expertise as one factor that qualified them to conduct audits. These included social statistics [P7], software engineering [P2], experience with HR [P4] or occupational psychology [P10]. In lieu of formally recognised standards for algorithm audit, our participants described drawing on internal and external codes of practice. For example, one organisation [P9] described aligning themselves with existing assurance standards, such as the ISAE 3000 issued by the International Auditing and Assurance Standards Board [2], as well as with draft audit criteria supplied by another organisation that participated in this study. Many of our interviewees reported conducting audits for companies they had pre-existing relationships with. In other auditing domains, such as financial audit, rules around the auditor-client relationship are tightly controlled [53]. Two participants felt the financial incentive in providing a bias audit under LL 144 might create challenges for independence [P1] [P12]. Several participants suggested that more flexibility about independence would increase the effectiveness of the bias audit: 'I needed [to work closely with] our clients to have better data governance practices for me to be able to do a better audit for them' [P14].

⁵For examples, see the earlier cited open-source repository, Gerchick & Watson (2023).

Others reported scepticism about the extent to which codes of practice for auditors under LL 144 appear to divert from existing practice elsewhere. For example, two participants [P10, P1] highlighted that the US Sarbanes Oxley Act 2002 strictly prohibits independent auditors from providing ‘certain non-audit services’ to clients [20]. Our interviews also revealed that some early LL 144 audits were being conducted in-house by employers on the grounds that another department within their organisation would qualify as independent. This touches on a tension between several interview subjects: some felt offering additional advisory services was not in opposition to the law’s definition of being an independent auditor [P3]. But others, like P5 and P6, felt a core component of being an independent auditor was that they could not provide advice or guidance on how to mitigate these issues. In their view, an auditor should be a one-off assessor of an AI system to provide “a review, a score, or a grade” that validates or certifies a company’s compliance with LL 144 [P6]. They also distinguished the role of an auditor as specifically assessing an AI product at a particular point in its lifecycle, as opposed to a continuous form of evaluation. This is why some auditors, like P5 and P6, explicitly rejected the term ‘auditor’ to describe their own services. From a business perspective, P5 and P6 also felt that offering companies the service of becoming ‘audit ready’ provided more opportunities to help their clients comply with other laws like the EU AI Act or the EU’s Digital Services Act. Like other elements of the law, what exactly constitutes a legitimate auditor remains vague and undefined within the law itself, other than the requirement of financial independence from the outcome of the audit report. Therefore, the responsibility to define auditor roles falls to employers and auditors themselves. Without clarity on the role of auditors, in practice, the power dynamics between existing actors in the field establish a working practice for implementation.

4.4 LL 144 created an opportunity for more companies to adopt wider responsible AI practices

‘This is your golden opportunity to change things in your organization. That’s how we try to put it.’ [P6]

Several of the participants we interviewed offered their clients additional services beyond those required under LL 144. P5, P6, P9, and P12, for example, all noted they offer wider responsible AI consultancy services, including support with data governance, wider AI ethics training, and even a SaaS platform to help clients enact responsible AI practices. Others, like P16, noted that they offer legal advice along with the audit. Some, like P1, P5, P6, P7, and P8, noted they are advising clients on how to get compliance-ready for other possibly forthcoming laws. Several participants noted that LL 144 created an incentive for companies and vendors to adopt wider responsible AI and data ethics practices, enabling participants to offer additional responsible AI services. P5 and P6 noted this kind of work was more lucrative than just the audit, which may be a one-off or once annual engagement. P14 noted that the law, while highly imperfect, had the effect of getting their clients to think about responsible AI and embed new language and ways of thinking into their corporate culture. According to P14, the audits ensured that

companies had a single AI system document and that their role as an auditor was to provide documentation and a ‘grounded language’ around responsible AI. The creation of a document enabled the wider staff to understand “that it’s absolutely appropriate for them to have a say in how AI is developed and applied in their company” [P5].

4.5 Auditors faced numerous obstacles around data practices, metrics and cultural attitudes toward data with vendors/clients

While technically straightforward, the audit process presents auditors with challenges concerning data practices and cultural attitudes in their engagements with vendors and clients. The question of how to procure the necessary data has become a common source of complexity for auditors. LL 144 only imposes obligations on the end-user—the employer—creating a relational dynamic between employer and vendor (with auditors mediating in the middle). This requires the end-user and the auditor to persuade the vendor for data access per the legal obligations.⁶ Even if the data is available, auditors encounter significant hurdles in accessing it due to vendor concerns around releasing proprietary systems or sensitive data. P14 noted that vendor cooperation was “key” and “one of the biggest obstacles to clients complying with the NYC law.” According to P14, “our only adversarial or third party was the vendor. So we had to take a very hard stance, getting our position and be kind of aggressive with the vendor to get us the data that we wanted so that we could produce that audit the way it’s meant to be.” P8 and P7 also referred to the role of non-disclosure agreements between clients and vendors as an obstacle that prevented access to the data necessary to do the audit. Often, auditors use tools like APIs or databases that allow controlled access to limited, but relevant, data. This tension between auditor and vendor not only extends the duration of audits but also underscores the importance of establishing streamlined processes for data sharing and accessibility.

Many participants cited a lack of guidance from DCWP here as a major contributing factor to data issues. As P14 noted, “the guidance still has a lot of discretion and flexibility,” and could be interpreted by different firms and vendors in different ways. This included guidance on how to select data for the audit; according to P14, “[the guidance] almost doesn’t talk about [sampling data]. So I mean, if we wanted to be in bad faith, we could look at our clients dataset and very easily pick rows that would give stellar results.” While P14 and other auditors acknowledged they would never do this, the ease of cherry-picking or biasing audit results, raising a serious risk to the legitimacy of the law. It also reduces the reliability of any measure of the success of the law that relies on tracking changes to impact ratios.

Many clients and vendors were initially resistant to acknowledging biases within their systems, making it difficult for auditors to help them understand their potential biases and potential mitigation strategies. P2 noted that some clients were “frugal” with their data but would provide full access after initial results indicated an

⁶The law is ambiguous about the extent to which an auditor may use a vendor’s own audit. ‘Pooled’ data from the vendor is allowed, as long as an employer also contributes their own data to the pool. See Raghavan et al. (2019) [46] on tradeoffs using pooled data in hiring algorithms

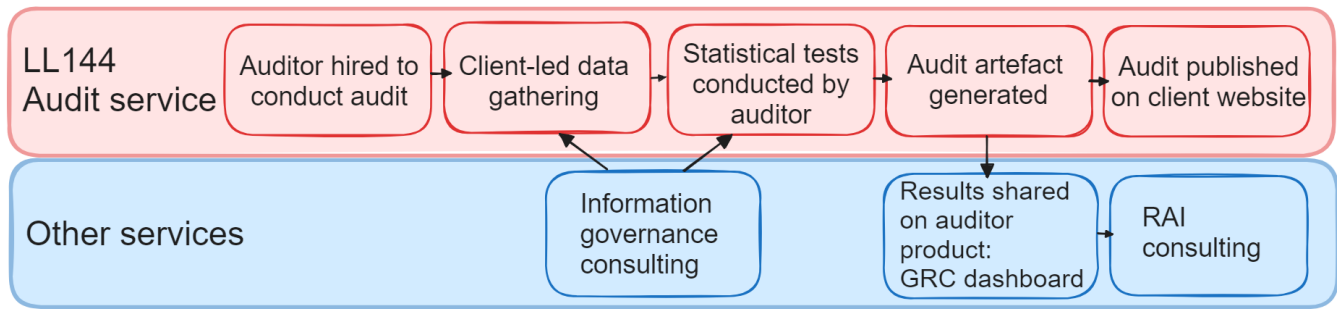


Figure 2: Services offered under the LL 44 regime by auditors

issue. P14 described how the process of conducting the audit was similar to “the stages of grief.” They reported stages of denial (“oh, no, we must have given you the wrong data or something like that”), despair (“oh, no, this is the worst thing in the world that could happen to us”), and finally acceptance from the client. This sense of denial matched the perspective of P8, who described clients who didn’t collect data on ethnicity or gender “because it’s sensitive” and then incorrectly claimed that this lack of data collection meant that bias was impossible.⁷ Bridging this gap between denial and acknowledgment proves to be a delicate and recurring challenge in the audit journey. The issue of ‘imputed or inferred’ demographic data adds another layer of complexity. While explicitly prohibited under LL 144 regulations, auditors noted that this practice persists, resulting in auditors often using names to infer demographic information like race, ethnicity, and gender. One auditor described a case where a client lacked any data on gender but made it up by assuming they had a 50/50 split because the national population in the USA is about a 50/50 male/female split. Instances of imputation/inferal are particularly problematic in diverse urban settings like NYC and raise questions about the validity of audits conducted under LL 144.

Additionally, the statistical test that LL 144 required auditors to use to determine disparate impact raised obstacles. P14 noted that it conflicts with ‘historical’ ways of doing disparate impact analyses where the protected group is pre-selected, before viewing the data. As a result, this caused the statistical validity of the test to go ‘out the window’ and enabled ‘cherry picking’ of data to tell a story that the system was not biased. Several auditors noted that having a mandatory, standardised metric for protected groups is important, but noted that a mandatory metric must be purpose-specific and appropriate [P6]. P10 noted the metric offered by DCWP changed part way through the implementation of the law from the average to the median, which caused delays and additional work.

5 DISCUSSION

Our interview findings and analysis of the law lead us to conclude that despite some modest successes in socialising bias audits among private companies, the law has fallen far short of its ambition to generate an effective bias audit regime that protects job applicants

⁷Federal regulations do require employers to keep anonymized demographic records of their hiring practices, recorded on the EE-01 form.

from discriminatory algorithmic tools. Of the thousands of employers in NYC, we could find only 19 who have publicly posted audits. While not a primary finding from this paper, we argue that the context of extremely low compliance is important for policymakers to note, both for making a judgement on the effectiveness of this auditing regime and for the wider implications for future AI regulatory initiatives. Many of the pinch points for algorithmic auditor practice identified in previous scholarship, including insufficient access, unclear metrics, and a lack of standardisation [8, 21, 41, 48, 56], hold true for auditor experience in LL 144, and so these findings are additive to existing calls for formalisation of practice. Below, we discuss four recommendations that policymakers seeking to develop similar algorithm auditing regimes should consider.

5.1 Auditing laws must have a theory of change that holds developers accountable, makes audits the source of action/enforceability and prioritises jobseekers

Councilwoman Laurie Cumbo has stated her motivations for sponsoring this law were to curb unjust hiring practices and promote equality [29]. In practice, however, the lack of audits and notices (and the poor readability/interpretability of those that do exist) means job candidates may not directly benefit. Our interviews do not reveal evidence of candidate outcomes being an influential driver or consideration behind the audits—according to P5, “the needs of the [candidate] are so different from the needs of an enterprise... and although we care about the [candidate], I don’t think that it is incumbent upon [us] to be able to satisfy the needs of that [candidate] at the same time as satisfying the needs of the enterprise.” This illustrates a tension between the personal motivation of some auditors to protect the public interest, and the market and regulatory conditions of algorithmic auditing as a service. At present, the LL 144 auditing regime is insufficient for creating accountability relationships between jobseekers and employers or vendors: the lack of prescribed mitigation steps in the event of adverse impact means a candidate would need to rely on plaintiff litigation for further action. Additionally, there is no requirement for the vendor to take action; an employer would need to convince vendors to improve the fairness outcomes of a discriminatory system. Consequently, the law forms a relational regime for industry actors to navigate –

employers who use AEDTs must navigate between their responsibility to LL 144, the potential censure of the EEOC, and the social cost of publishing a non-favourable bias audit. To remedy this, we recommend future auditing laws take an ecosystem view to accountability, understanding that there is ‘distributed responsibility’ [14] between vendors, users/deployers and other actors. In practice, this might involve placing audit requirements on vendors or hiring platforms (we note that this was complicated in the NYC context due to the city’s jurisdictional limits). One of our participants proposed that vendors should be obliged to publish the audits of all employers using their tools. We join Raji et al. (2022) and others in arguing for centralised reporting: for future regimes, it is critical that a regulatory body hosts an audit repository, which will assist candidates in making comparisons between different employers and law firms and regulators in enforcement. Future laws could mandate the disclosure of additional context about the employer and the tool (including contact details), which could make the audit results easier for researchers and the public to parse and scrutinise. Audits could also be supplemented by other kinds of responsible AI practices like algorithmic impact assessments that seek to address accountability issues in different parts of the product development lifecycle [27, 35, 51]. Lastly, to avoid the over-representation of industry voices in public consultations around auditing regimes, it is critical for policymakers to intentionally engage with stakeholders who will be impacted by these systems.

5.2 Auditing laws need clear and meaningful metrics and definitions of key terms

According to our participants, a major source of NYC LL 144’s failure is the last-minute addition of language around AEDTs that ‘substantially assist’, which drastically reduced the number of companies in scope. Additionally, the statistical metric used by the NYC LL 144 law (disparate impact, based on the four-fifths convention) is heavily rooted in US law and is not applicable to other jurisdictions outside of the USA. Policymakers creating similar regimes must ensure they carefully consider the terms and metrics used and ensure that these align with their local cultural context. Terms should always be defined in a way that is weighted in favour of protections for those impacted by these technologies over those building them. Regulators or departments responsible for overseeing these regimes can take steps to issue updated guidance about the use of these tools. Crucially, policymakers must also not shy away from enforcement, which is a primary driver of compliance in these regimes.

5.3 Establishing an auditing ecosystem requires more standards for and oversight of auditors

Many of the observed challenges surrounding the regime are about limited evaluation procedures and oversight and enforcement capacity. Auditors are largely left to their own devices in defining ‘good’ practice in this space, and interviewees suggested that the DCWP may be ill-equipped to enforce the law. Further, without a pool of certified or accredited auditors, under this regime, any organisation without a financial stake in the AI system or audited company can conduct audits. Even in good faith, this leaves the

door open for inconsistencies in practice and prevents oversight that might enable ecosystem-level learning.

Our interviews have also highlighted that algorithm auditing regimes can lead to a range of actors delivering various services beyond the strict definition of an ‘independent auditor.’ These include companies providing the service of making companies ‘audit ready’ or advising them on steps they can take to improve their data and responsible AI practices. We join Birhane et al. (2024) in calling for the recognition of the intellectual diversity and range of competencies this array of organisations can bring to audit [8]. Our interviews suggest these roles are essential, but they must be clearly defined and differentiated in an auditing regime. Future auditing regimes need to establish clear requirements and standards for what constitutes an ‘independent’ auditor. A next step for national governments and/or standards bodies would be to define standards of practice for algorithm auditors, for the benefit of auditors and regulators⁸. These standards should also set out methodological details surrounding the approach to auditing for second and third-party audits. This could build on, for example, the UK Information Commissioner’s Office guidance for AI auditing [40]. As suggested elsewhere [32, 48], the financial auditing model could also provide a steer: for example, an algorithm auditing equivalent to the US Public Company Accounting Oversight Board [44] could be established. We also draw attention to emerging non-governmental bodies with a remit to define algorithm audit standards, such as the newly established International Association of Algorithm Auditors [1]. Previous work on AI accountability and assurance has identified a need for an ‘AI ombudsman’ [30] as an authority to field complaints and grievances. We propose that a similar model could be adopted to mediate AI auditors. Additionally, increased regulatory capacity at the local and national levels is required for regulators to enforce audit compliance.

5.4 Auditing laws must enable smoother data collection for performing tests and providing access to data for certified auditors

Our findings underscore the need for legal frameworks that facilitate seamless data collection, empowering auditors to conduct meaningful assessments. One prevailing approach relies on auditors fostering positive relationships with a system’s developers, emphasising concerns that evaluations often hinge on selective disclosure, potentially favouring auditors who portray a positive image of the system. The issue of access is further complicated by a lack of trust in companies releasing only favourable data, a challenge exacerbated by the fact that APIs are often developed and managed by large, for-profit entities. In recent years, companies providing tools like APIs, intended for public system data use, have raised costs and restricted tool usage. Notably, platforms such as Reddit have introduced charges for API access, while others have shut down their APIs altogether. Consequently, auditors frequently find themselves obligated to declare their intent for data usage, introducing uncertainty and potential bias into the auditing process.

⁸Models of such guidance exist in other domains and which regulations might easily reference. For example, the ISAE 3000 series of standards for assurance engagements defines the roles and duties of third-party auditors generally.

From a legal perspective, auditors face legal vulnerability due to existing language within privacy and anti-hacking laws. For instance, the US Computer Fraud and Abuse Act criminalises ‘unauthorised access to systems, particularly regarding data collection without permission [39]. The ambiguity surrounding methods classified as authorised versus unauthorised creates uncertainty when auditors use tools like data scrapers to access system data, potentially violating platform Terms of Service. In response to these challenges, we advocate for legal clarity and protection for algorithm auditors to conduct thorough and meaningful evaluations. Precision in legal language, including terms of use for data collection methods and tools, is essential to ensuring that auditors are shielded from unwarranted legal repercussions.

6 CONCLUSION

In this paper, we explore the implementation LL 144, the first law creating an algorithm bias audit regime. While an admirable first attempt, its vague definitions and narrow scope ultimately failed to drive a robust implementation of bias audits, leaving the real work of auditing to be decided by industry actors, auditors, and hiring platforms. The law created a warped incentive structure resulting in under-compliance for employers in scope, but also over-compliance from vendors and platform not in scope. This is in part a consequence of placing all obligations on the software’s end user—the employer using the AEDT—despite the end-user typically not building or even having access to the aspects of the system typically required for a robust algorithmic audit. In addition, the law grants the end-user total discretion over whether their system is in scope, and offers multiple loopholes for employers to move out of scope without addressing discriminatory outcomes. While the law specifies what the audit should include, it fails to specify the role of the auditor, creating confusion about the appropriate credentials and practices of independent auditors. Consequently, there is little standardisation in LL 144 audit services. The law has driven a secondary market of responsible AI services that, in many cases, are becoming the primary activity under this regime. Auditors are challenged by navigating the accountability relationships created by the law, including negotiating system access from developers and with legal counsel about whether a system is in scope. Clarifying definitions and the responsibilities between vendors, employers, and jobseekers may help establish clearer accountability relationships and mitigate potential non-compliance. We join calls from across the FAcCT community that similar laws should address the capacity building of algorithm audit expertise, standards of practice, and robust oversight mechanisms [48]. For mandated audits to become a meaningful accountability mechanism in AI, it is critical that these regimes apportion accountability in a way that delivers beneficial outcomes for people impacted by these systems.

ACKNOWLEDGMENTS

This project was supported in part by the Open Society Foundation and Siegel Family Endowment. The paper benefited significantly from a related project conducted in collaboration with researchers from the Cornell University Citizens and Technology Lab, lead by J. Nathan Matias. The authors would like to thank Catherine Gregory, Sohaib Malik, Rebecca Ghani, Octavia Reeve, and Ranjit Singh

for their thoughtful comments and revisions to this paper. This work was supported by the grant sponsor UKRI, under the grant number EP/V00784X/1 (Trustworthy Autonomous Systems (TAS) Hub). This work was supported in part by the National Science Foundation Award #1704425.

7 ETHICS STATEMENT

7.1 Ethics consideration statement

The research protocol was reviewed by an ethics review committee at the Ada Lovelace Institute and determined to be low risk. The specific conditions of the study were not subject to any legally-mandated independent research ethics review. However, care was taken to adhere to standard research ethics norms in the construction of the study, conducting the interviews, and handling human subject data. All interview participants signed an informed consent and confidentiality agreement form agreeing to voluntarily participate in this study (with the option to withdraw at any time). Participants were issued a document outlining research objectives and the purpose/expectations of the interview, as well as specific stipulations around how interview data would be stored, retained, and governed for the duration of the study. The data will inform a separate policy report outside of this FAcCT paper, for which participants gave their informed consent. As part of our interviews, we asked participants to provide their name and their organisation, which count as personally identifiable data categories under the UK General Data Protection Regulation (UK GDPR) and the US Privacy Act of 1974. The two organisations leading this study signed a cross-border data governance agreement to steward data responsibly. Interview data was anonymised at the point of transcription by the lead interviewer, based in the UK, to ensure personal data remained in the jurisdiction of collection, per UK GDPR. Interview transcripts are stored on a password-protected cloud storage system. Due to the disclosed steps, we do not foresee significant risk or impact for participants as a result of participating in this study.

7.2 Researcher positionality statement

This research was conducted as part of a partnership between two civil society research institutes based in the United States and the United Kingdom. Both research institutes are independent, do not receive funding or stipends from governments or the technology industry, and thus take full research and editorial control and accountability for this paper and the findings therein. Our site of study was a NYC regulatory regime aimed at targeting NYC-based company practices, with the primary beneficiaries being NYC jobseekers. Expertise and local knowledge of, for example, federal employment law was therefore beneficial for research legitimacy in this context. Four authors are North American and one is British. Three authors have educational backgrounds in social science disciplines, and two have backgrounds in the intersection of computer science and social science. The authors adopt a sociotechnical lens for AI, which understands that AI systems are derived from both technical components and social and societal histories, norms, and values [30]. We firmly acknowledge the context for the genesis of this law, which was sponsored by Black Council Women in accordance with the principles of (algorithmic) justice and equity, and offer our express support for these principles, which also underpin

the authors' work in the domain of AI and society. We also want to acknowledge and offer gratitude for the contributions of the actors and organisations, including researchers, journalists, and activists, who have advanced both the intellectual debate and methodological rigour of third-party algorithmic auditing as a method for AI accountability.

7.3 Adverse impact statement

The authors do not foresee a significant risk of adverse or unintended impact on the participants surrounding the publication of this work.

REFERENCES

- [1] [n. d.]. IAAA - International Algorithmic Auditors Association. ([n. d.]). <https://iaaa-algorithmicauditors.org/>
- [2] 2013. ISAE 3000 Revised for IAASB. *International Auditing and Assurance Standards Board* (2013). <https://www.iaasb.org/publications/international-standard-audit-engagements-iaae-3000-revised-audit-engagements-other-audits-or>
- [3] 2023. Comments Received by the Department of Consumer and Worker Protection on Proposed Rules related to Automated Employment Decision Tools. *New York City Consumer and Worker Protection* (2023). <https://www.nyc.gov/assets/dca/downloads/pdf/about/PublicComments-Proposed-Rules-Related-to-Automated-Employment-Decision-Tools-Updated.pdf>
- [4] 2024. Code of Federal Regulations. *29 CFR Part 1607 - General Principles* (2024). <https://www.wired.com/story/opinion-new-york-city-surveillance-battle-offers-national-lessons/>
- [5] Ifeoma Ajunwa, Sorelle Friedler, Carlos E Scheidegger, and Suresh Venkatasubramanian. 2016. Hiring by algorithm: predicting and preventing disparate impact. *Available at SSRN* (2016).
- [6] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *California law review* (2016), 671–732.
- [7] 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.
- [8] Abeba Birhane, Ryan Steed, Victor Ojewale, Briana Vecchione, and Inioluwa Deborah Raji. 2024. AI auditing: The broken bus on the road to AI accountability. *arXiv preprint arXiv:2401.14462* (2024).
- [9] Jenny Brennan, Lara Groves, Elliot Jones, and Andrew Strait. [n. d.]. AI Assurance? <https://www.adalovelaceinstitute.org/report/risks-ai-systems/>
- [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] Albert Fox Cahn. 2021. New York City's Surveillance Battle Offers National Lessons. *Wired* (2021). <https://www.wired.com/story/opinion-new-york-city-surveillance-battle-offers-national-lessons/>
- [12] Pew Research Center. 2023. AI in Hiring and Evaluation of Workers: What People Think. (2023). <https://www.pewresearch.org/internet/2023/04/20/ai-in-hiring-and-evaluating-workers-what-americans-think/>
- [13] Kathy Charmaz. 2023. *Constructing grounded theory: A practical guide through qualitative analysis*. sage.
- [14] Jennifer Cobbe, Michael Veale, and Jatinder Singh. 2023. Understanding accountability in algorithmic supply chains. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1186–1197.
- [15] European Commission. 2024. The Digital Services Act package | Shaping Europe's digital future. (2024). <https://digital-strategy.ec.europa.eu/en/policies/digital-services-act-package>
- [16] Equal Employment Opportunity Commission. 2007. Employment Tests and Selection Procedures. (2007). <https://www.eeoc.gov/laws/guidance/employment-tests-and-selection-procedures>
- [17] Equal Employment Opportunity Commission. 2022. Definitions of Race and Ethnicity Categories. *Data Collection* (2022).
- [18] Equal Employment Opportunity Commission. 2023. Assessing Adverse Impact in Software, Algorithms, and Artificial Intelligence Used in Employment Selection Procedures Under Title VII of the Civil Rights Act of 1964. (2023). <https://www.eeoc.gov/laws/guidance/select-issues-assessing-adverse-impact-software-algorithms-and-artificial>
- [19] Australian Competition and Consumer Commission. 2020. Trivago misled consumers about hotel room rates. (2020). <https://www.accc.gov.au/media-release/trivago-misled-consumers-about-hotel-room-rates>
- [20] US Congress. 2002. The Sarbanes Oxley Act. (2002). <https://sarbans-oxley-act.com/>
- [21] Sasha Costanza-Chock, Inioluwa Deborah Raji, and Joy Buolamwini. 2022. Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1571–1583.
- [22] New York City Council. 2021. A Local Law to amend the administrative code of the city of New York, in relation to automated employment decision tools. (2021). <https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=4344524&GUID=B051915D-A9AC-451E-81F8-6596032FA3F9&Options=Advanced&Search=>
- [23] Donald R Deis Jr and Gary A Giroux. 1992. Determinants of audit quality in the public sector. *Accounting review* (1992), 462–479.
- [24] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 259–268.
- [25] Marissa Gerchick and Brooke Watson. 2023. Tracking Automated Employment Decision Tool Bias Audits. (2023). <https://github.com/aclu-national/tracking-ll144-bias-audits>
- [26] Ellen P Goodman and Julia Trehu. 2022. Algorithmic Auditing: Chasing AI Accountability. *Santa Clara High Tech. LJ* 39 (2022), 289.
- [27] Lara Groves, Jenny Brennan, Inioluwa Deborah Raji, Aidan Peppin, and Strait. 2022. Algorithmic impact assessment: a case study in healthcare. (2022). <https://www.adalovelaceinstitute.org/wp-content/uploads/2022/02/Algorithmic-impact-assessment-a-case-study-in-healthcare.pdf>
- [28] Indeed. 2023. The Indeed Global AI Survey: Your Guide to the Future of Hiring. (2023). <https://www.indeed.com/lead/the-indeed-ai-report?hl=en&form>
- [29] Irina Ivanova. 2020. New York City wants to restrict artificial intelligence in hiring. *CBS News* (2020). <https://www.cbsnews.com/news/new-york-city-artificial-intelligence-hiring-restriction/>
- [30] Elliot Jones, Jenny Brennan, Connor Dunlop, and Andrew Strait. 2023. Keeping an eye on AI. *Ada Lovelace Institute* (2023). https://www.adalovelaceinstitute.org/wp-content/uploads/2023/09/ALI_Keeping-an-eye-on-AI-2023.pdf
- [31] Lauren Kirchner. 2023. New York City Moves to Create Accountability for Algorithms. *ProPublica* (2023). <https://www.propublica.org/article/new-york-city-moves-to-create-accountability-for-algorithms>
- [32] Adriano Koshiyama, Emre Kazim, Philip Treleaven, Pete Rai, Lukasz Szpruch, Giles Pavey, Ghazi Ahamat, Franziska Leutner, Randy Goebel, Andrew Knight, et al. 2021. Towards algorithm auditing: a survey on managing legal, ethical and technological risks of AI, ML and associated algorithms. (2021).
- [33] David Leslie. 2019. Understanding artificial intelligence ethics and safety. *arXiv preprint arXiv:1906.05684* (2019).
- [34] Steve Lohr. 2023. A Hiring Law Blazes a Path for A.I. Regulation. *The New York Times* (2023). <https://www.nytimes.com/2023/05/25/technology/ai-hiring-law-new-york.html>
- [35] 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).
- [36] New York City Department of Consumer and Worker Protections. 2023. Automated Employment Decision Tools: Frequently Asked Questions. (2023). <https://www.nyc.gov/assets/dca/downloads/pdf/about/DCWP-AEDT-FAQ.pdf>
- [37] New York City Department of Consumer and Worker Protections. 2023. Automated Employment Decision Tools (Updated). (2023). <https://rules.cityofnewyork.us/rule/automated-employment-decision-tools-updated/>
- [38] Colorado Division of Insurance. 2023. Regulation 10-1-1 Governance and Risk Management Framework Requirements for Life Insurers' Use of External Consumer Data and Information Sources, Algorithms, and Predictive Models. (2023). https://drive.google.com/file/d/1dPKJcDo76iHfJZDopQEhTDCmKbuYnNI/view?usp=embed_facebook
- [39] US Department of Justice. 2015. 9-48.000 - Computer Fraud and Abuse Act. (2015). <https://www.justice.gov/jm/jm-9-48000-computer-fraud>
- [40] Information Commissioner's Office. 2020. Guidance on the AI auditing framework Draft guidance for consultation. Information Commissioner's Office. (2020). <https://ico.org.uk/media/2617219/guidance-on-the-ai-auditing-framework-draft-for-consultation.pdf>
- [41] Victor Ojewale, Ryan Steed, Briana Vecchione, Abeba Birhane, and Inioluwa Deborah Raji. 2024. Towards AI Accountability Infrastructure: Gaps and Opportunities in AI Audit Tooling. *arXiv preprint arXiv:2402.17861* (2024).
- [42] UK Parliament. 2022. Online Safety Act 2023. (2022). <https://bills.parliament.uk/bills/3137>
- [43] Jules Pattison-Gordon. 2023. Colorado Aims to Prevent AI-Driven Discrimination in Insurance. *GovTech* (2023). <https://www.govtech.com/policy/colorado-aims-to-prevent-ai-driven-discrimination-in-insurance>
- [44] PCAOB. [n. d.]. Driving improvement in audit quality to protect investors. ([n. d.]). <https://pcaobus.org/>
- [45] Evani Radiya-Dixit and Gina Neff. 2023. A Sociotechnical Audit: Assessing Police Use of Facial Recognition. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1334–1346.
- [46] 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.
- [47] 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.
 - [48] 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*. 557–571.
 - [49] Algemene Rekenkamer. 2022. An Audit of 9 Algorithms used by the Dutch Government - Report - Netherlands Court of Audit. (2022). <https://english.rekenkamer.nl/publications/reports/2022/05/18/an-audit-of-9-algorithms-used-by-the-dutch-government>
 - [50] Yvette D. [D-NY-9] Rep. Clarke. 2022. Algorithmic Accountability Act of 2022. (2022). <https://www.congress.gov/bill/117th-congress/house-bill/6580/text>
 - [51] Andrew D Selbst. 2017. Disparate impact in big data policing. *Ga. L. Rev.* 52 (2017), 109.
 - [52] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*. 59–68.
 - [53] SHRM. 2022. Automation and AI in HR. (2022). https://advocacy.shrm.org/SHRM-2022-Automation-AI-Research.pdf?_ga=2.112869508.1029738808.1666019592-61357574.1655121608
 - [54] Nopmanee Tepalagul and Ling Lin. 2015. Auditor independence and audit quality: A literature review. *Journal of Accounting, Auditing & Finance* 30, 1 (2015), 101–121.
 - [55] Elizabeth Anne Watkins, Michael McKenna, and Jiahao Chen. 2022. The four-fifths rule is not disparate impact: a woeful tale of epistemic trespassing in algorithmic fairness. *arXiv preprint arXiv:2202.09519* (2022).
 - [56] 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.
 - [57] Lucas Wright, Roxana Mika Muenster, Briana Vecchione, Tianyao Qu, Pika Cai, Alan Smith, COMM/INFO 2450 Student Investigators, Jacob Metcalf, and J. Nathan Matias. 2024. Null Compliance: NYC Local Law 144 and the Challenges of AI-Algorithm Accountability. *ACM Conference on Fairness, Accountability and Transparency* (2024).

A APPENDICES

A.1 Example audit report, prepared by ConductorAI for National Broadcasting Company (NBC)

Audit Results

Sex Categories			
	<u># Of Applicants</u>	<u>Scoring Rate</u>	<u>Impact Ratio</u>
Male	2,063,618	47%	0.99
Female	2,007,072	47%	1

Race/Ethnicity Categories			
	<u># Of Applicants</u>	<u>Scoring Rate</u>	<u>Impact Ratio</u>
Hispanic or Latino	470,904	45%	0.89
White	1,457,444	46%	0.9
Black or African American	796,447	48%	0.96
Native Hawaiian or Pacific Islander	14,904	45%	0.9
Asian	941,554	50%	1.00
Native American or Alaska Native	29,612	43%	0.85
Two or More Races	216,700	46%	0.92

Intersectional Categories					
	<u>Sex</u>	<u>Race/Ethnicity Categories</u>	<u># Of Applicants</u>	<u>Scoring Rate</u>	<u>Impact Ratio</u>
Hispanic or Latino	Male		231,334	44%	0.87
Hispanic or Latino	Female		237,315	45%	0.89
Non-Hispanic or Latino	Male	White	725,976	46%	0.9
Non-Hispanic or Latino	Male	Black or African American	351,245	48%	0.94
Non-Hispanic or Latino	Male	Native Hawaiian or Pacific Islander	7,656	45%	0.89
Non-Hispanic or Latino	Male	Asian	547,776	51%	1
Non-Hispanic or Latino	Male	Native American or Alaska Native	17,000	41%	0.82
Non-Hispanic or Latino	Male	Two or More Races	99,271	46%	0.91
Non-Hispanic or Latino	Female	White	723,288	46%	0.91
Non-Hispanic or Latino	Female	Black or African American	442,836	49%	0.97
Non-Hispanic or Latino	Female	Native Hawaiian or Pacific Islander	7,117	46%	0.90
Non-Hispanic or Latino	Female	Asian	391,030	50%	0.99
Non-Hispanic or Latino	Female	Native American or Alaska Native	12,262	44%	0.88
Non-Hispanic or Latino	Female	Two or More Races	113,712	47%	0.93

Figure 3: Example audit report, prepared by ConductorAI for NBC

A.2 Table 2: Participant ID

Participant ID	Type of organization	Background of interviewee	Audit service
P1	Non-profit	Entrepreneur, economics	Audit certifier
P2	Data science firm	Entrepreneur, software engineering	Audit provider
P3	Law firm	Labour economics, litigation	Audit provider
P4	Self-employed consultant	HR, AI governance	N/A
P5	GRC platform*	Data science	Pre-audit provider
P6	GRC platform*	Technology policy	Pre-audit provider
P7	Data science firm*	Quantitative social science	Audit provider
P8	Data science firm*	Data science	Audit provider
P9	Consultancy/technology development firm	Ethics	Audit provider, additional services provider
P10	GRC platform	AI governance, management, occupational psychology	Audit provider, additional services provider
P11	Consultancy/audit firm*	AI ethics, governance	Audit provider, audit certifier
P12	Consultancy/audit firm*	Data science	Audit provider, audit certifier
P13	Consultancy/software firm	Data science	Audit provider, additional services provider
P14	GRC platform	Data science, mathematics, law, public policy	Audit provider, additional services provider
P15	Self-employed consultant, previously of a talent assessment platform	Cognitive science, AI ethics & governance	AEDT platform developer
P16	Law firm	Law, data science	Audit provider, additional services provider
P17	Employment platform	Data science, computational social science	N/A

Table 2: Participant ID

Entries marked with an asterisk represent instances where we interviewed two representatives from the same organisation.