

Null Compliance: NYC Local Law 144 and the Challenges of Algorithm Accountability

Lucas Wright*
law323@cornell.edu
Cornell University
Ithaca, NY, USA

Roxana Mika Muenster
rm858@cornell.edu
Cornell University
Ithaca, NY, USA

Briana Vecchione
briana@datasociety.net
Data & Society Research Institute
New York City, NY, USA

Tianyao Qu
tq34@cornell.edu
Cornell University
Ithaca, NY, USA

Pika (Senhuang) Cai
sc3322@cornell.edu
Cornell University
Ithaca, NY, USA

Alan Smith
alan.smith@consumer.org
Consumer Reports
Yonkers, NY, USA

COMM/INFO 2450 Student
Investigators[†]
Cornell University
Ithaca, NY, USA

Jacob Metcalf
jake.metcalf@datasociety.net
Data & Society Research Institute
New York City, NY, USA

J. Nathan Matias
nathan.matias@cornell.edu
Citizens and Technology Lab, Cornell
University
Ithaca, NY, USA

ABSTRACT

In July 2023, New York City became the first jurisdiction globally to mandate bias audits for commercial algorithmic systems, specifically for automated employment decisions systems (AEDTs) used in hiring and promotion. Local Law 144 (LL 144) requires AEDTs to be independently audited annually for race and gender bias, and the audit report must be publicly posted. Additionally, employers are obligated to post a transparency notice with the job listing. In this study, 155 student investigators recorded 391 employers' compliance with LL 144 and the user experience for prospective job applicants. Among these employers, 18 posted audit reports and 13 posted transparency notices. These rates could potentially be explained by a significant limitation in the accountability mechanisms enacted by LL 144. Since the law grants employers substantial discretion over whether their system is in scope of the law, a null result cannot be said to indicate non-compliance, a condition we call "null compliance." Employer discretion may also explain our finding that nearly all audits reported an impact factor over 0.8, a rule of thumb often used in employment discrimination cases. We also find that the benefit of LL 144 to ordinary job seekers is limited due to shortcomings in accessibility and usability. Our findings offer important lessons for policy-makers as they consider regulating algorithmic systems, particularly the degree of discretion to grant to regulated parties and the limitations of relying on transparency and end-user accountability.

*Please direct all correspondence to Lucas Wright at law323@cornell.edu

[†]See Appendix A.4 for full list of student investigators.



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/10.1145/3630106.3658998>

CCS CONCEPTS

- **Social and professional topics** → **Governmental regulations**;
- **Computing methodologies** → **Artificial intelligence**.

KEYWORDS

algorithm audit, compliance, transparency, bias

ACM Reference Format:

Lucas Wright, Roxana Mika Muenster, Briana Vecchione, Tianyao Qu, Pika (Senhuang) 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 Algorithm Accountability. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FACT '24)*, June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3630106.3658998>

1 INTRODUCTION

On the cusp of potential major changes to the AI regulation landscape in many jurisdictions, New York City implemented the world's first law mandating the conducting and publishing of algorithmic bias audits for commercial products in July 2023. NYC Local Law 144 (LL 144) mandates that any NYC-based private employer or City agency that deploys certain automated employment decision tools (AEDTs) in the hiring or promotion process must conduct a disparate impact study¹ for race and gender features, and make the **audit report** available to the general public via their website (Appendix A.1). The employer must also provide a **transparency notice** to any job seeker, informing them about the use of an AEDT (Appendix A.2).

This law includes elements of interest to any government seeking to regulate algorithmic systems. For example, the third-party audit

¹LL 144 specifically names these audits as 'bias audits,' but because they only measure a specific type of bias (disparate impact) against a narrow set of protected demographic classes (race and gender) we choose to use the term 'disparate impact audit/study'. Please see Appendix A and Section 2.4 for a fuller description of the meaning of 'disparate impact', how that is measured, and its relation to both LL 144 and the so-called 'four-fifth's rule.'

requirement creates a market for auditors, following practices in accounting [64], pollution [59], and compliance monitoring [34]. The requirement of notice follows customs in employment law [37]. Yet the outcomes of these policy components have yet to be established in real-world algorithm governance.

This paper presents early analysis of the publicly-available outcomes of this historically-important algorithmic governance regime, two years after its passage and five months after it came into effect. Collaborating with 155 student investigators acting as model job seekers, we report what can be learned about employer compliance with the law, report qualitative and quantitative findings on the job seeker experience, and analyze the contents of published audit reports. Given the reported widespread usage of AEDTs and large workforce in NYC, we found surprisingly low rates of affirmative compliance with the required public audit reporting (5%) and transparency notices (3%) in our sample. Yet the accountability structures of LL 144 create a conundrum for investigators: the high level of discretion granted to employers to decide if their systems are in scope means that any null result cannot be considered non-compliance. Research can measure a compliance rate, but not a non-compliance rate. To make sense of this situation, we introduce the term **null compliance** to describe a situation in which non-compliance cannot be ascertained.² Null compliance does not name a third state between compliance and non-compliance that the employer may hold; rather, it describes the state of public knowledge *about* compliance within this algorithmic transparency regime. Although this may be the first named example of null compliance in algorithmic governance, we anticipate seeing further examples in other algorithmic accountability regimes that are primarily driven by strategies of transparency and consumer choice.

This law has created market opportunities for algorithm auditors and increased publicly-available information about some AEDTs used by NYC-based employers. Yet we find that the discretion afforded to employers by the law hinders the full potential utility of mandated transparency. Null compliance makes the extent of AEDT use by employers unclear, blunts the effectiveness of transparency requirements, undermines the choice structures intended to protect job seekers, and obfuscates any evidence about changes or reductions in discrimination over time. Taken together, these problems in a promising and landmark law make it impossible for researchers to know if LL 144 is reducing employment discrimination, increasing it, or making it harder to reliably monitor overall.

2 REGULATING AUTOMATED EMPLOYMENT DECISION TOOLS IN NYC WITH LOCAL LAW 144

LL 144 grew out of longstanding but troubled efforts to introduce a degree of transparency to the use of algorithmic systems by the municipality of New York City [4, 35, 36]. These efforts shifted toward regulating commercial systems with potential civil rights impacts when in early 2020 City Commissioner Laurie Cumbo sponsored a

²Null compliance' has been used in other fields, such as accounting, biomedical research, and instrument calibration to describe the extreme end of non-compliance, meaning zero compliance, where compliance with a standard can be measured on a scale. Our usage is distinct: in transparency-driven algorithm regulation where compliance is a binary state, null compliance describes not being able to ascertain the state.

statute requiring employers to disclose and audit machine learning systems used to assist in hiring [17]. The law was enacted in late 2021, and it assigned the Department of Consumer and Workers Protection (DCWP) responsibility for rule-making and enforcement. The DCWP went through two rounds of rule-making and public comment—including two delays in the implementation date—before issuing final rules in April of 2023, with implementation in July of 2023 [49]. The law covers employers that are located in NYC, including remote jobs that primarily report to a NYC office.

2.1 Implementation Details of LL 144

The final version of LL 144 imposes two related but distinct obligations upon employers³, who are typically end-users of these systems.⁴ First, employers must hire an independent audit provider (described as an expert with no financial stake in the company or the outcome of the audit) to conduct an annual disparate impact audit of the system and post it publicly on their website. This **audit report** must focus on features of race/ethnicity, gender, and the intersection of both as defined by the U.S. Equal Employment Opportunity Commission (EEOC). Second, the employer must provide job seekers with a **transparency notice** that their application will be analyzed by an AEDT, and provide job seekers with an opt-out mechanism to request a human review process if one is available or otherwise required by law.⁵ The transparency notice must either accompany the job posting or be sent to applicants via correspondence. No law requires disclosure of bias among human reviewers, so job seekers do not have information about the alternative option.

The law specifies civil penalties for violators, ranging from \$500 to \$1500 per violation per day.⁶ Enforcement relies on complaints; the law does not provide the DCWP with proactive investigatory or discovery powers, nor does it grant a private right of action to job seekers. To the best of the authors' knowledge, no complaints have been submitted to the DCWP at the point of publication. Due to the regulatee discretion detailed below, job seekers are unlikely to know whether any AEDT they may encounter during the application and interview should properly be in scope, and researchers and regulators cannot determine if the absence of an audit indicates absence of an AEDT.

2.2 Sources of Discretion Offered to Employers

The drafters of LL 144 and the DCWP made several choices about the accountability structures of the law that significantly shape the conditions of this study, particularly the discretion granted to the employers that use AEDTs. **First**, LL 144 directly **imposes obligations upon employers only**, although most employers that use AEDTs lease them from vendors and/or recruitment platforms that train and maintain the models [49]. As a municipality, NYC

³The law also applies to employment agencies, which procure and sometimes hire candidates for employers. We found no audits from these agencies.

⁴If the AEDT is built in-house the employer may be both end-user and developer. We were unable to confirm any such case in our survey.

⁵The law specifies that employers are *not* obligated to provide an alternative/human-only evaluation method *unless* alternative methods are required by law, i.e., for medical or disability accommodations. Nonetheless, the theory of change behind LL 144 emphasizes job seeker agency to not apply to jobs with discriminatory, or any, AEDTs, and thus a form of opt-out is always available.

⁶The text of the law does not specify if the violation is per system, per applicant, or per job posting, making the cost of non-compliance unclear.

has regulatory powers over employers in its jurisdiction but not developers and platforms that are domiciled elsewhere. Additionally, employment anti-discrimination law focuses on moments of decision-making or evaluation between employer and employee, for which software platforms are typically not directly responsible.⁷

Second, the DCWP establishes **no central public repository** for audit reports, relying solely on employers to post their own. **Third**, the law **does not bar systems** that cause disparate impact [49]—it only requires reporting the rates. This choice places LL 144 in a complicated relationship to federal anti-discrimination law that does set a floor via the so-called ‘four-fifths rule’ (see Appendices 1 and 2, and section 2.4 for more detail).

Given these features, the near-total discretion given to employers to interpret the scope of LL 144 has important consequences for the success of the law. It is common practice to grant significant discretion to the regulatee to self-determine if and how the scope of a rule applies to them based on principles promulgated by the regulatory agency [57]. While LL 144 is not outside of typical parameters of regulatee discretion, it does appear to create a tension between discretion and transparency that ultimately blunts the effectiveness of the law, a result we name ‘**null compliance**’.

The DCWP defines an AEDT as software that uses machine learning techniques to “substantially assist or replace discretionary decision-making” to mean [49]:

- i. to rely solely on a simplified output (score, tag, classification, ranking, etc.), with no other factors considered; or
- ii. to use a simplified output as one of a set of criteria where the simplified output is weighted more than any other criterion in the set; or
- iii. to use a simplified output to overrule conclusions derived from other factors, including human decision-making

This is an awkward mix of technical features and organizational features that differs from technical definitions of machine learning [1], introducing interpretational complexity and granting employers wide latitude over whether their systems are in scope. Given criteria ii. and iii., two employers that use identical models could grant those models different degrees of influence over their respective hiring decisions—an entirely organizational matter—and thereby have different regulatory statuses. Likewise, a job seeker could encounter the same vendor’s AEDT during applications to two different employers and be protected by an audit at one but not the other, entirely contingent on how the employer describes their internal processes. Given those conditions, it is reasonable to expect that many employers and their legal advisors would be either confused about, or seek ways to legally avoid, the scope of the law.

While there is no reliable source about what percentage of employers utilize AEDTs of any type, what surveys do exist show that AI systems for recruitment, hiring, and HR are widespread and rapidly growing in reach [31, 62]. Vendors offer a wide variety of AI services to hiring managers, from matching candidates with openings to providing summaries of resumes to scheduling interviews;

⁷This introduces significant challenges for an algorithmic accountability regime, however, because typically neither the end-user nor their hired auditor has direct access to the backend functionality and training data of the developer/vendor that are needed for algorithmic audits. Another paper using data from this study examines this issue in depth (see REDACTED, forthcoming).

LL 144 applies only to AEDTs that are used at any point in the process to facilitate decisions regarding advancement/screening out of candidates [49]. While ethically-complex use cases grab headlines (e.g., personality/psychological profiles, cognitive tests, and computer vision tools) lower-stakes AEDTs appear to be more common. Although LL 144 is built with them in mind, it does not appear that any vendors currently offer tools that fully replace human discretion in the hiring process with complete automation; thus any attempt to track the effect of the law would need to focus on systems that “substantially assist” humans.

2.3 Local Law 144 and Anti-Discrimination Law

Finally, the relationship between LL 144 and federal anti-discrimination law is critically important to the outcomes of this study and the conundrum of null compliance. LL 144 requires an audit, but is silent on the results of that audit. The so-called ‘four-fifths rule’ [3, 50, 69] that emerged from US federal anti-discrimination regulations and jurisprudence looms large in employment discrimination and algorithmic fairness (see **Appendix** for key terms and historical details). The four-fifths rule establishes a baseline against which impermissible discriminatory outcomes (aka ‘disparate impact’) can be identified by comparing relative rates of selection between demographic groups. If the less-selected group’s selection rate falls below four-fifths (80%) of the baseline then there is a presumption that disparate impact may be occurring and warrants further investigation, typically through a complaint to the EEOC. The four-fifths rule is not genuinely a rule or law; it is best described as a rule of thumb or convention that is ingrained in hiring practices, and employers and developers are highly incentivized to reduce potential scrutiny by targeting that threshold. Nonetheless, LL 144 does not require that systems meet any discrimination threshold, including the four-fifths rule. Nor does LL 144 provide any guidance for remediation of systems when audits disclose disparate impact. The DCWP’s FAQ notes that other laws on employment discrimination still apply but does not cross-reference them in any way [49].⁸

In other words, the most relevant statistical criterion for judging discrimination in employment is consequential to the first algorithm auditing law *primarily by its absence*. This situation is central to understanding the conundrum of null compliance: one regulator (the DCWP) has demanded transparency about an activity that another regulator (the EEOC) has jurisdiction over but would not be able to observe in the usual course of business. The EEOC would typically need to pursue a discrimination complaint to find this information. Nor has any federal body established guidance about safe harbors for audits conducted in compliance with a local or state law. Such safe harbors would protect employers from liability for voluntarily contributing to transparency regimes. Thus, employers complying with LL 144 by posting audit reports might also be opening themselves to other types of liability, creating conflicting regulatory pressures [43].

That context speaks to LL 144’s theory of change, a term used by policy scholars to explain the link between policy details and

⁸The first case of algorithmic hiring bias—for age discrimination—was settled with the EEOC in 2023, but it was for a system that automatically screened out all older applicants and not a typical disparate impact scenario [15, 26].

its desired ends[5, 6]. LL 144 requires transparent auditing so job seekers are empowered to make informed choices and employers and vendors are incentivized to deploy fair(er) models, but it does not set a floor for permissible rates of discrimination. It also creates a market for algorithm auditing services, and a sandbox for experimenting with auditing techniques. The outcomes of LL 144 need to be measured against that theory of change.

3 THE MEANING OF NULL COMPLIANCE: NOT THE SAME AS NON-COMPLIANCE

Because LL 144 gives employers so much discretion, any attempt to study compliance with the law will face difficulties estimating compliance rates due to *presumed under-compliance and voluntary over-compliance*. The field of design differentiates between designers and end-users of systems who are not involved in their making [47]. The law places accountability only upon the end-user of AEDT systems (the employers) to procure and post an audit, grants them near-total discretion over whether their system is in scope, and offers them many chances to move out of scope [48]. The law also offers no formal mechanism for challenging these employer decisions. Consequently, when investigators, regulators, and job seekers cannot find an audit report or transparency notice, they cannot call it non-compliance.

To explain this situation, we introduce the term: ‘**null compliance**’. Null compliance describes a state in which the absence of evidence of compliance cannot be ascertained as non-compliance because the investigator lacks the information to determine if the regulated party’s actions or products are in scope of the regulation. Null compliance is a state of public knowledge that emerges from a transparency regime, not a state between compliance and non-compliance that the regulatee holds. In this situation, researchers can disclose an affirmative compliance rate (i.e., a regulatee affirms that they are in scope and compliant), but they cannot disclose a non-compliance rate. We anticipate seeing further examples in other algorithmic accountability regimes that emphasize transparency and yet offer discretion to end-users.

In this paper, we report on information about whether researchers have observed an audit report and/or transparency notice on an employer’s public web page. We use the term ‘null compliance’ to refer to employers where our study was unable to verify compliance and also unable to verify non-compliance. Null compliance should not be taken to indicate non-compliance in this paper—absent legal investigatory powers that we lack, *we cannot say that any given case of null compliance is actually non-compliance*. A finding of a null compliance in our report may occur due to any combination of the following conditions made possible by LL 144’s particular accountability arrangements:

- (1) The employer does not deploy an AEDT.
- (2) The employer deploys AEDTs, but is not aware of LL 144.
- (3) The employer deploys AEDTs, but has self-determined that their system is outside of the scope of the law.
- (4) The employer deploys AEDTs, has self-determined that the law applies to their system, and is still searching for or waiting on a third-party auditor, which is a new industry.

- (5) The employer deploys AEDTs, has conducted an audit, and has decided that they should not post the results publicly as a risk management strategy.
- (6) The employer deploys AEDTs, has conducted an audit, and complies with the law by providing job seekers with the transparency notice over email or the postal service rather than a public website.
- (7) The employer deploys an AEDT and has conducted an audit, but has posted either/both the audit report and/or transparency notice in hard-to-discover location(s).

4 PRIOR SCHOLARSHIP ON MEASURING COMPLIANCE AND STUDYING OPT-OUT DECISIONS

In this section, we summarize prior scholarship on measuring regulatory compliance and informed choice.

4.1 Measuring Compliance with Transparency Regulations

Researchers in other fields have also wrestled with the difficulty of working with evidence produced through monitoring and transparency laws. Because these policies are designed to influence regulated behavior while aiding enforcement [59], they create incentives that complicate the study of regulated behavior. These incentives pressure firms toward under-compliance and over-compliance alike. The resulting uncertainty leaves society unclear about the outcomes of regulation. Our definition of null compliance incorporates an understanding of these sources of uncertainty.

4.1.1 Deliberate Ignorance. Especially when research might expose firms to legal risks, “trapped administrators” can conclude that their self-interest depends on preventing or evading research [8]. In environmental policy and public health, scientists have described monitoring as a “co-evolutionary race” between researchers who seek knowledge and firms that seek to evade that research [19]. The threat to research validity from deliberate ignorance is so common that it is now the subject of a social science subfield on “agnotology” [51].

For researchers studying LL 144, deliberate ignorance could lead to under-counts in base rates of AEDT usage as well as biased estimates of racial and gender disparities in AEDTs. In the case of LL 144, employers have strong financial incentives to avoid conducting and publishing a bias audit. While the cost of non-compliance with LL 144 is a theoretical maximum of \$547,500 per year, civil action by federal regulators or private cases brought by litigators can be much larger [15, 26]. Since non-compliance with LL 144 prevents third parties from knowing details that could expose employers to liability, they have a strong incentive toward deliberate ignorance.

4.1.2 Information Access Control. When regulated entities do choose to know things, they can still manage regulatory risk by limiting access to that information. Economists have described how self-reporting policies incentivize firms to publish results only when the costs and uncertainty of self-reporting are lower than those of withholding information [34]. Access control, which economists describe in terms of information asymmetries, can create challenges

for researchers seeking to estimate rates of compliance [59]. Computer scientists have described a similar dynamic where firms are incentivized to share data with researchers if they believe the results of research will be favorable [72]. For LL 144, access control could lead to an under-count of AEDTs, with mostly favorable audit reports published.

4.1.3 Mis-measurement and Fabrication. Mis-measurement by regulated entities is a further threat to the validity of research on regulatory compliance. When regulations impose penalties on the basis of measurement, firms may respond by reporting false or adjusted information. In the study of air quality, for example, researchers have found systematic evidence of fabricated and false data from regulated firms and third-party auditors [21, 39, 61]. Scientists have also found evidence of systematic bias in official air quality measurements in cases where local politicians have discretion in their placement [18]. In the case of this study, we are unable to determine whether mismeasurement is occurring.

4.1.4 Creative Behavior at Decision Rules. Researchers have long observed that measurements become less reliable the more they are used for decision-making, especially when observations are close to a decision rule [9, 27]. For example, education researchers wrestle with the reliability of measurements that are used for decisions such as school allocation, graduation, and college admissions [45, 55]. Parallel scholarship has studied the problem of publication bias in science that result from creative efforts by scientists to navigate p-value thresholds [32].

LL 144 includes two decision-rules that invite creative behavior. The four-fifths rule (Appendix A.1) is often used as a rule of thumb in employment discrimination law. Since companies have discretion over whether their hiring software is governed by the law, some companies might use their discretion under LL 144 to avoid EEOC investigation.

4.1.5 Over-Compliance. Researchers seeking to accurately measure regulated behaviors also face the problem of over-compliance. Without over-compliance, researchers could claim that compliance rates are under-counts, but over-compliance adds uncertainty in the opposite direction.

When regulators create a market for compliance monitoring, multiple actors have incentives toward over-compliance by following policies that may not apply to them or by complying to a degree that is not required [60]. Risk-averse firms can over-comply out of an abundance of caution, even when a law does not apply to them or when the scope of a recent law remains uncertain [7]. Firms sometimes comply with monitoring and reporting regimes to grow their reputation compared to competitors [22]. Finally, third party monitors market their commercial services to firms where auditing services are not legally required. Any of these incentives could lead some firms to over-comply with LL 144. While over-compliance might generate public goods by giving job seekers more choices and increasing the transparency of hiring software, it also undermines the accuracy of research on the use of AEDT systems and compliance with the law.

4.1.6 Obfuscation and Unusability. Even when firms comply with expectations that they warn people about potential risks from their products [20], they exercise considerable agency over the visibility

and accessibility of disclosures [38, 65]. Obfuscation affects research when disclosures are made so inaccessible that researchers cannot find them.

Since LL 144 gives job seekers the right to request a costly alternative to the AEDT, employers have strong financial incentives to make disclosures difficult to find. Employers have limited guidance and no case law to inform the user experience. Consequently, researchers seeking out these disclosures will likely under-count them.

4.2 Notice and Consent in Other Areas of Regulation

Promoters of disclosure and consent policies argue that they enable individuals to make informed decisions in cases of significant information asymmetry. These policies enable the state to regulate firms indirectly by balancing those asymmetries [67]. For this to work, people need to be sufficiently informed and able to make choices.

Research in behavioral science and computing has found that, at best, disclosure and consent policies only help some people under some circumstances and that they can backfire [30, 52]. In employment law, behavioral scientists have found that equal opportunity and diversity statements can reduce applications from minority job seekers [37]. Such statements can also increase discrimination when they convince minority job-applicants to disclose information that subsequently activates biased hiring decisions [33]. The disclose-and-opt-out model of LL 144 is also similar to many data privacy regulations. In this model, individuals are expected to read detailed information about a company's practices and make an informed decision. Computer scientists have observed that this model requires hundreds of hours per year per person [42]. Even with time, people still struggle to make choices in their own interests [65, 66].

4.3 Characteristics of Successful Notice and Consent

Notice and consent depends on informed decision-making, which relies on the accessibility, comprehensibility, and usability of information. As many researchers have observed, regulation can create strong incentives for firms to undermine each of these characteristics.

First, information needs to be accessible. If job seekers cannot find transparency notices or bias audit reports, then they effectively have no choice, even if one is available in some technical sense. In other areas of technology policy, firms have tended to make disclosures inaccessible, requiring a "scavenger hunt" just to exercise one's rights [29].

Second, disclosures need to be comprehensible enough for a job seeker to make an informed choice [73]. Ongoing literature in bioethics, psychology, and design is exploring exactly what it means to make an informed decision [56, 71, 74, 75]. At minimum, an informed decision is one in which people are able to reason about the possible consequences based on reliable information about the possible choices [71].

Finally, the success of notice and consent policies depends on usability. A growing toolbox of "dark patterns" offers people the illusion of choice while steering people toward decisions that benefit

the designers to their own detriment [10, 41]. These dark patterns are especially common in areas where firms and the public have competing incentives, including price discrimination and regulatory compliance [40].

5 METHODS

In this paper, we set out to summarize how employers are publishing information to job seekers, describe the user experience they are creating for job seekers, and analyze the contents of the audit reports. To do so, we organized 155 student investigators to record public information they found on the websites of 391 employers in a manner modeled after job-seeking. We then validated the information, invited employers to correct any errors, and analyzed the resulting corpus. We have published our archive as an open resource.⁹

To inform the design of the full research protocol, we collected examples of bias audit reports a week after the law went into effect on July 5, 2023. By one month after the law came into effect, we had found 19 audit reports, 14 from employers plus 5 from software vendors not covered by LL 144. This pilot dataset was largely identical to a crowd-sourced dataset later assembled by civil society actors [25].

Our employer sample included 568 employers that had hired graduates of the classes of 2021 and 2022, based on data from the university careers office.¹⁰ Employers were included if they hired a communication or information science major or hired at least two students from any major. We also included employers identified in our pilot study and employers that had been ranked as the top 100 internship providers in the country.¹¹ We excluded businesses in states well beyond NYC or no presence in the US for a total of 511. Our final list of 391 employers was a random sample from this list.

5.1 Modeling job seekers on Employer Websites

Computer scientists studying regulatory compliance and user experience have typically taken two kinds of approaches. In one approach, researchers use software to crawl websites [40]. Other researchers take the approach we used by organizing investigators to undertake large number of searches that model user behavior [29].

We recruited 155 student investigators through an extra credit opportunity in a large lecture class for Communication and Information Science students at Cornell University in New York State, where many graduates are hired by employers in NYC¹². In an hour-long session, we trained them to search for audit reports and transparency notices as if they were a job seeker. They searched the company's website and profiles on LinkedIn and Indeed for no more than 30 minutes per employer to model the amount of time a dedicated job seeker might reasonably invest into the process.

⁹<https://citizensandtech.org/2024-algorithm-transparency-law>

¹⁰This sample is limited to the employers that hire Cornell students, which tend to be credentialed/professional roles. AEDT usage and transparency may be different in other sectors of the job market.

¹¹The full list of these employers can be found at <https://www.nationalinternday.com/2023-top-internship-programs>

¹²Participants in this exercise who consented to sharing their names are acknowledged in Appendix A.4

In order to protect against adverse consequences for the students, they did not submit applications.¹³

For each employer, investigators recorded whether the employer offered jobs in New York City at the time of their search. They also recorded evidence of any bias audit report and transparency notice they found, including screenshots, links, document files, and a description of the section of the website where they found the material. To minimize false negatives, we encouraged investigators to record any and all material that might plausibly be related to LL 144.

Student investigators collected data in two stages from October 24 through November 9, 2023. In the first stage, each investigator was randomly assigned one employer to research, assigning each company to two investigators. In the second stage, we assigned each investigator three more employers to research, with less overlap.

5.2 Data Verification and Aggregation

We next produced a list of verified reports and transparency notices from student investigator findings. Because we instructed investigators to be as inclusive as possible to minimize false negatives, they sometimes recorded other types of disclosures or statements of rights as being related to AEDTs. We reviewed the evidence they submitted with their results, checked the website if necessary, and removed false positives.

We also identified cases where employers using the same AEDT vendor published duplicate audit reports. In our analysis of employers, we count each employer that provided an audit report. When analyzing the audit reports themselves, we included each unique audit report once. In cases where an audit report reported impact ratios for multiple performance goals, we counted impact ratios separately since they report unique results.

5.3 Survey of Employers

We also contacted employers with two rounds of emails from Consumer Reports to further verify whether our research team might have missed public information about LL 144. This yielded two additional audit reports, out of 26 responses (Appendix A.3). False negatives in this study are not necessarily or solely an indication of researcher error: they may also illustrate the significant challenge of discovering these regulatory artifacts for an actual job seeker.

6 FINDINGS

Based on the data collected from employers, we offer findings on observed compliance, the user experience created by employers, and the AEDT impact ratios published in audit reports.

6.1 Summarizing Observed Compliance with Local Law 144

Among 267 employers posting open jobs in NYC, we found 14 audit reports and 12 transparency notices (Table 1). Our sample included 124 employers that did not have a job opening in NYC at the time of the search. Of those companies, we found 4 bias audits and 1 transparency notice; such an outcome was expected because the

¹³This method could not observe transparency notices that were only supplied after someone applies.

transparency notices typically appear in conjunction with an active listing. Notably, only 11 employers published both an audit report and a transparency notice that met our understanding of the law's requirements, suggesting incomplete compliance among employers that have determined their use of AEDTs falls within the scope of the law.

While this study can count instances of compliance with LL 144, the problem of null compliance prevents third parties such as ourselves or job seekers from verifying employers' compliance with the law in a manner that advances the laws' transparency goals. Null compliance describes the phenomenon of not being able to interpret the absence of an audit report and/or transparency notice in our findings, which could result from at least seven possible causes that each have different legal implications (see Section 4). Given this, the absence of an audit report or transparency notice is not an indicator of non-compliance on the part of any individual employer. Because the law affords a level of discretion to employers that makes the denominator impossible to know, we likewise cannot make claims about the percentage of employers in compliance with the law. Furthermore, while we can describe the audit reports that have been made public, the denominator problem prevents us from making general claims about AEDT systems in New York City, including any trends in algorithmic bias that might result. This distortion from null compliance in public knowledge about AEDTs is a direct result of the design of the accountability mechanism in LL 144.

In follow-up communications, employers explained how they exercised the discretion afforded by the law. Of the 26 firms that responded to our survey, 23 reported that LL 144 did not apply to them and 3 reported that LL 144 applied to them in some other way. Of the 3 employers, 2 supplied a document that met our definition of a bias audit. One company reported that while they do use AEDT systems more widely, they have chosen not to use these systems for NYC-based jobs in order to comply with LL 144.

6.2 The User Experience of Compliance with LL 144

We report on the user experience of compliance with LL 144 on the basis of quantitative and qualitative findings from student investigators and an analysis of the materials that employers published. We consider three overall questions: the user experience of LL 144 compliance, the readability of bias reports, and the user experience of opting out.

6.2.1 The Inaccessibility of Transparency Notices and Audit Reports. The theory of change behind LL 144 assumes that employers will make information accessible to job seekers who can make informed decisions. In this study, trained student investigators assessed whether the a motivated job applicant could locate the information that LL 144 mandates. Overwhelmingly, students could not do so, describing the experience as challenging, time-consuming, and frustrating.

Student investigators repeatedly reported the difficulty of locating the transparency notice and bias report on employer websites. Often, they were unable to locate transparency notices or bias audit reports within the 30-minute allotted search time, despite clicking through 17 pages and 19 pages on average, respectively. Employers

that did publish the information placed it in very different parts of their websites. Disclosures sometimes appeared in the footer of the website or in the careers FAQ; other times, it was accessible through a drop-down menu or downloadable as a PDF. Given the challenges faced by trained student investigators, we expect that job seekers will only rarely encounter audit reports or transparency notices even when available.

6.2.2 Ambiguity Amidst Other Disclosures. Information about LL 144 was often placed among a larger list of other legal disclaimers and notices in confusing ways. Consequently, investigators struggled to identify whether a given notice fulfilled the requirements of LL 144 or some other regulation. Notices sometimes included unnecessary information such as ambiguous labels or vague mission statements. One student investigator reported: "I spent a lot of time sifting through information on the website, but I could not find it anywhere, despite the company claiming they value diversity, equity, and inclusion and minimizing unconscious bias."

For student investigators, this cacophony was amplified by many similar documents and disclosures, especially on large employers' websites: Privacy notices, financial audit reports, diversity statements and reports, or financial compensation bias audit reports were all mixed in AEDT transparency notices or bias audit reports. These reports were difficult to distinguish from each other even by trained investigators. Of records that student investigators thought might be audit reports 31% were audit reports and 20% were transparency notices as defined by LL 144. We therefore do not believe most job seekers could distinguish these statements under normal circumstances.

6.2.3 Ambiguity Over the Applicability of LL 144. Some employers demonstrated compliance while also claiming that the law did not apply. One perplexing example presented an AEDT bias audit report but prefaced it with an extensive disclaimer that, in the employers' opinion, their tool did not actually constitute an AEDT as defined by LL 144.

In several cases, what the student investigators identified as a transparency notice was actually the opposite. Instead of disclosing the use of AEDTs, these statements informed job seekers that the company did not use AEDTs when hiring or that they do not solely rely on them in hiring decisions (which is not enough to exclude them from the law's scope).

6.2.4 The Usability of Audit Reports and Transparency Notices. Once job seekers find audit reports, they need sufficient, clear information to make informed decisions. Our study identified 18 employer-provided bias audit reports, with some providing the same report as their AEDT vendor, leaving us with 13 unique reports.¹⁴

None of these reports included enough explanation to inform job seekers or guide a decision to opt-out of applying. Audit reports were typically a PDF document with information about the audit provider, the employer, and the AEDT vendor, along with data tables (see Appendix A.1). They usually included the scoring rate,

¹⁴The final rule permits employers to pool data from multiple employers' data that use the same vendors' model, as long as the employer also contributes their data to the pooled dataset. Yet since audit reports cannot and should not be the same between employers, they are required to publish unique reports.

	Total	Audit Report (N)	Audit Report (%)	Transparency Notices (N)	Transparency Notices (%)
Employers listing NYC Jobs	267	14	5%	12	4%
Employers not listing NYC Jobs	124	4	3%	1	1%
All employers	391	18	5%	13	3%

Table 1: Across 391 employers, researchers found 18 audit reports and 13 transparency notices

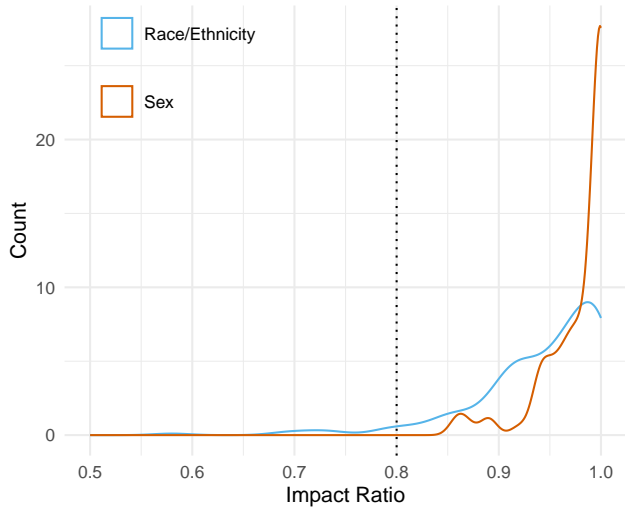


Figure 1: Kernel density showing the distribution of impact ratios for sex and race/ethnicity categories (N = 13 audit reports and 386 impact ratios) among publicly-posted audit reports

selection rate, and impact ratio across sex and race/ethnicity, with minimal explanation. Only three transparency notices mentioned the types of data collected and used in automated decision-making.

6.3 Impact Ratios in Published Bias Audit Reports

Most published audit reports reported impact ratios that were above an 0.8 threshold that would provide a presumption of non-discriminatory outcomes under federal law. While being on either side of the 0.8 threshold is not considered dispositive of either discrimination or non-discrimination in any jurisdiction, any numerical evidence of a disparate impact could draw negative attention and be used adversely against the employer in litigation (A.1, A.2). As stated above, the ‘four-fifths rule’ is correctly understood as a convention that is highly incentivized. Out of 386 total measurements, nine impact ratios were below 0.8, while all others reported impact ratios above this threshold.

We can only speculate on the reasons why 96% of published audit reports included impact ratios above the 0.8 “threshold”. It’s possible that employers are more likely to share the results of favorable audits or that the estimates are mismeasured. It is also possible, though we think it is unlikely, that almost no employers in our sample used AEDT systems with performance below the 0.8 threshold. In a companion study to this one [28], audit industry

workers indicated in interviews that, in their experience, many, if not the majority, of AEDTs on the market violate the four-fifths rule. Multiple audit service providers further indicated that they had clients who paid for an audit for an in-scope system but declined to post it publicly.

We did observe one employer that may be withholding an audit report. Since this employer unpublished their audit report before our study began, we did not include it. The report, which others have archived [25], shows that the audit provider reported AEDT impact ratios below 0.8 in two demographic areas. We cannot explain why the report was unpublished—it is possible they ceased using AEDTs in NYC. When we contacted the employer, they did not respond.

Why would any organization pay for an audit and then fail to comply with the law’s public posting requirement? Because LL 144 does not set a floor for acceptable impact ratios, employer compliance with LL 144 could make them vulnerable to enforcement from the EEOC or private litigation. Legal counsel may advise them that non-compliance with LL 144 is less risky than providing evidence for such litigation. While our data cannot be used to prove such publication bias in audit reports, our evidence is fully consistent with what we would expect in that situation and what auditors in Grove et al. [28] indicated is the case.

The audit reports regularly excluded data for some race/ethnicity subgroups, offering the argument that the subgroup represented less than 2% of the data being used for the audit, which is in line with the DCWP’s rules adopted to address concerns about distorted results due to low statistical power. Across all bias audit reports analyzed, employers omitted information for 20% of impact ratios for this reason [24]. This is important to consider given that most omitted impact ratios fell under already marginalized indigenous groups. In one case, impact ratios were not reported at all: Holistic AI’s bias audit of Lendlease provided no results for impact ratios across race/ethnicity.

7 DISCUSSION

The theory of change behind LL 144 is that transparency, the risk of EEOC enforcement, and job seeker agency will drive employers to adopt less biased systems and thereby create a market for well-governed and fair(er) AEDTs, or disincentivize use of such systems altogether. Policy-makers can learn important lessons about how LL 144’s reliance on transparency and sustaining a market for auditing service does and does not make that possible. Ultimately, the effectiveness of the law should be judged on one outcome: reduced discriminatory outcomes in hiring and promotion where automated systems are deployed. Due to the problem of null compliance, the law and its current implementation do not enable anyone to know

if that outcome is being achieved within individual employers or across the job market.

7.1 Null Compliance in LL 144 Hinders Enforcement and Transparency

Overall, the problem of null compliance hinders the law's function and prevents researchers from knowing if the law's purposes are being achieved. Because the law gives employers significant discretion (3), absence of information is null compliance rather than non-compliance, creating problems for enforcement and research alike. Null compliance creates strong barriers to enforcement. In such cases, investigators or complainant job seekers cannot tell whether the law does not apply, whether an employer is refusing to conduct a bias audit, or whether they are hiding an unfavorable audit. Null compliance also prevents researchers from compiling general information or claims about algorithmic hiring decision systems. The distribution of published impact ratios gives us reason to believe that publication bias could be causing a significant over-estimate of the fairness of AEDTs used to hire people in New York City.

Within audit reports, null compliance can obscure employment disparities for minority communities due to provisions in the law that permit employers to exclude data from audit reports. Since employers are in practice excused from reporting impact ratios for groups that they have hired the least, LL 144 is unable to achieve transparency or accountability for those groups. This sample size issue has important implications for algorithm accountability laws to address the rights of American Indians, Alaska Natives, and other indigenous groups.

The most significant change that could be made to LL 144 or subsequent laws would be to remove much of the discretion granted to the regulated party, which in this case is the end-user. In our judgment, the most effective route to this would be to attach the scope of the law to the purpose of the system. As it stands, employers can off-ramp from the regulatory decision tree by claiming (correctly or incorrectly) that their decision-making process does not 'substantially' rely upon the outputs of the AEDT, or by using techniques that evade the technical definition of AEDTs. While this off-ramp may serve the purpose of incentivizing employers to reduce their reliance on AEDTs in decision making, it also creates the counter-intuitive outcome that a job seeker may encounter the same system across different employers but only be protected by LL 144 in some cases. Regulations could instead define AEDTs to include any system that outputs a simplified score or ranking from aggregated job seeker data, regardless of how the employer claims it is used or how it was built. This would reduce the regulatees' discretion (and confusion) over whether their AEDT is the right kind of AEDT to be regulated. If employers were required to publish consistently-formatted audit reports to a common repository, overall transparency would more easily be achieved.

7.2 Confusing User Experiences and Partial Reporting Requirements Undermine job seeker Agency

In theory, LL 144 empowers job seekers to make informed decisions about whether to accept an AEDT based on transparency notices

and audit reports. Yet as the law was implemented by employers, job seekers cannot be reasonably expected to access, understand, or make use of that information. Transparency notices were hard to find and buried amidst many other disclosures. If job seekers do find the appropriate information, the legal and technical jargon in reports would prevent anyone except legal and technical experts from making an informed decision. Impact ratios and their meaning in anti-discrimination law are not simple to understand. Audit reports are not obligated to explain even basic concepts, such as that a lower score means a more discriminatory outcome.

A more effective regulation would provide clear, consistent guidance to employers on implementing notice and opt-out, including an opt-out for all job seekers. In other areas, federal regulators are considering 'click to cancel' policies offering guidance on opt-out designs [63]. Even if employers published easy-to-understand audit reports with a simple user experience, job seekers still could not make truly informed decisions about AEDTs in the absence of information about the alternative. Looking at an impact ratio, the job seeker might rightly worry whether choosing human reviewers might expose them to even more bias than choosing AEDT. The most effective transparency regulation would require employers to report on the non-AEDT alternative as well.

7.3 Could LL 144 Reduce Employment Discrimination?

In theory, publicly-available audit reports could provide evidence of changing impact ratios over time, guiding the evaluation of the overall outcomes of LL 144 on employment discrimination. Due to problems of null compliance and the lack of a mandated central repository for audit reports, tracking this is prohibitive. An improved law could set a floor for permissible impact ratios, even if that floor is fundamentally arbitrary and blunt. LL 144 does not set such a floor and instead relies on the 'four-fifths rule' implicitly; it makes no reference to the rule but requires auditors to use exactly the equations and disclose results that any disparate impact investigation would make use of. This creates perverse incentives that make it less risky for companies to withhold data rather than prevent or reduce discrimination.

Absent some sort of safe harbor, any employer that uses a system with an impact ratio below 0.8 will need to decide if complying with LL 144's transparency requirements will provide information for EEOC enforcement or private litigation. And because an impact ratio above 0.8 is not a guaranteed shield against litigation, publishing even apparently good results could conceivably feed adverse litigation outcomes. They may reasonably judge that it is highly unlikely that piecemeal enforcement by a local jurisdiction will be more costly than a federal civil suit.

8 CONCLUSION

Algorithmic employment tools are notoriously unreliable and ill-defined, imposing significant harm on workers [46, 53, 54]. Policy experts have argued that the chief outcomes of disclosure-based regulations are the capacities they build to achieve at least some compliance [58] and facilitate grass-roots contestation [44]. Despite its significant flaws, Local Law 144 is achieving these first steps. In our judgment, the route by which LL 144 is most likely to drive

change in the employment algorithm industry is by forcing the vendors' customers to pause, measure and record outcomes, discuss internally about what their tools actually do, reflect on how they use them, and deliberate on whether their AEDTs are even necessary. We have some evidence that this is happening, from employers who claim to have stopped using these tools in New York City. If fewer employers use hiring algorithms, the rate of algorithmic bias will definitionally go down, even though it is hard to know if that would increase or decrease employment discrimination overall.

ACKNOWLEDGMENTS

We are deeply grateful to the students of COMM/INFO 2450 for their thorough and extensive collaboration on this project and to the Cornell University Department of Communication for supporting this classroom research. J. Nathan Matias received financial support from the Center for Advanced Study in the Behavioral Sciences, the Lenore Annenberg and Wallis Annenberg foundation and the Siegel Family Endowment. Researchers at the Data & Society Research Institute were supported by a grant from Omidyar Network, and in part by a grant from the Open Society Foundations. We thank our other collaborators on related projects, namely Alexandra Mateescu (who supported the pilot audit collection), Ranjit Singh, Brian Chen, Lara Groves, Andrew Strait, and Alayna Kennedy. We also thank the Data & Society RAW Materials Seminar participants for their insightful comments. And we are grateful to the anonymous peer reviewers and conference committee members for supporting this venue for scholarship.

REFERENCES

- [1] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2023. *Fairness and machine learning: Limitations and opportunities*. MIT Press.
- [2] Solon Barocas and Andrew D. Selbst. 2016. Big data's disparate impact. *California law review* (2016), 671–732. <https://www.jstor.org/stable/24758720> Publisher: JSTOR.
- [3] Philip Bobko and Philip L Roth. 2004. THE FOUR-FIFTHS RULE FOR ASSESSING ADVERSE IMPACT: AN ARITHMETIC, INTUITIVE, AND LOGICAL ANALYSIS OF THE RULE AND IMPLICATIONS FOR FUTURE RESEARCH AND PRACTICE. In *Research in Personnel and Human Resources Management*. Research in Personnel and Human Resources Management, Vol. 23. Emerald Group Publishing Limited, 177–198. [https://doi.org/10.1016/S0742-7301\(04\)23004-3](https://doi.org/10.1016/S0742-7301(04)23004-3)
- [4] Albert Fox Cahn. 2019. The first effort to regulate AI was a spectacular failure. *Fast Company* (Nov. 2019). <https://www.fastcompany.com/90436012/the-first-effort-to-regulate-ai-was-a-spectacular-failure>
- [5] Paul Cairney. 2016. *The politics of evidence-based policy making*. Springer.
- [6] Paul Cairney. 2019. *Understanding public policy: theories and issues*. Vol. 2. Bloomsbury Publishing.
- [7] John E. Calfee and Richard Craswell. 1984. Some effects of uncertainty on compliance with legal standards. *Va. L. Rev.* 70 (1984), 965. https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/valr70§ion=41 Publisher: HeinOnline.
- [8] Donald T. Campbell. 1969. Reforms as experiments. *American psychologist* 24, 4 (1969), 409. <https://psycnet.apa.org/record/1969-17253-001> Publisher: American Psychological Association.
- [9] Donald T. Campbell. 1979. Assessing the impact of planned social change. *Evaluation and program planning* 2, 1 (1979), 67–90. <https://www.sciencedirect.com/science/article/pii/014971897990048X> Publisher: Elsevier.
- [10] Evan Caragay, Katherine Xiong, Jonathan Zong, and Daniel Jackson. 2023. Beyond Dark Patterns: A Concept-Based Framework for Software Design Ethics. <http://arxiv.org/abs/2310.02432> arXiv:2310.02432 [cs].
- [11] Rohit Chopra, Kristen Clarke, Charlotte A. Burrows, and Lina M. Kahn. 2023. Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems. https://www.ftc.gov/system/files/ftc_gov/pdf/EEOC-CRT-FTC-CFPB-AI-Joint-Statement%28final%29.pdf
- [12] Equal Employment Opportunity Commission. 2007. Employment Tests and Selection Procedures. <https://www.eeoc.gov/laws/guidance/employment-tests-and-selection-procedures>
- [13] Equal Employment Opportunity Commission. 2022. The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees. <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence>
- [14] 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. <https://www.eeoc.gov/laws/guidance/select-issues-assessing-adverse-impact-software-algorithms-and-artificial>
- [15] Equal Employment Opportunity Commission. 2023. iTutorGroup to Pay \$365,000 to Settle EEOC Discriminatory Hiring Suit. <https://www.eeoc.gov/newsroom/itutorgroup-pay-365000-settle-eeoc-discriminatory-hiring-suit>
- [16] Equal Employment Opportunity Commission. 2023. Navigating Employment Discrimination in AI and Automated Systems: A New Civil Rights Frontier - Transcript. <https://www.eeoc.gov/meetings/meeting-january-31-2023-navigating-employment-discrimination-ai-and-automated-systems-new/transcript>
- [17] 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. <https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=4344524&GUID=B051915D-A9AC-451E-81F8-6596032FA3F9&Options=Advanced&Search=>
- [18] Priyanka deSouza and Patrick L. Kinney. 2021. On the distribution of low-cost PM2.5 sensors in the US: demographic and air quality associations. *Journal of Exposure Science & Environmental Epidemiology* 31, 3 (May 2021), 514–524. <https://doi.org/10.1038/s41370-021-00328-2> Number: 3 Publisher: Nature Publishing Group.
- [19] Thomas Dietz, Elinor Ostrom, and Paul C. Stern. 2003. The struggle to govern the commons. *science* 302, 5652 (2003), 1907–1912. <http://science.sciencemag.org/content/302/5652/1907.short>
- [20] Hardy Cross Dillard and Harris Hart. 1955. Product Liability: Directions for Use and the Duty to Warn. *Virginia Law Review* (1955), 145–182. <https://www.jstor.org/stable/1070193> Publisher: JSTOR.
- [21] 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. <https://academic.oup.com/qje/article-abstract/128/4/1499/1850465> Publisher: MIT Press.
- [22] Stefanie Engel. 2006. Overcompliance, labeling, and lobbying: The case of crede goods. *Environmental Modeling & Assessment* 11 (2006), 115–130. Publisher: Springer.
- [23] Gene Foreman, Daniel R. Biddle, Emilie Lounsbury, and Richard G. Jones. 2022. *The ethical journalist: Making responsible decisions in the digital age*. John Wiley & Sons.
- [24] S. Michael Gaddis (Ed.). 2018. *Audit Studies: Behind the Scenes with Theory, Method, and Nuance*. Springer International Publishing, Cham. <https://doi.org/10.1007/978-3-319-71153-9>
- [25] Marissa Gerchick and Brooke Watson. 2023. Tracking Automated Employment Decision Tool Bias Audits. <https://github.com/aclu-national/tracking-ll14-bias-audits> original-date: 2023-08-09T19:46:06Z.
- [26] Annelise Gilbert. 2023. EEOC Settles First-of-Its-Kind AI Bias in Hiring Lawsuit. *Bloomberg Law* (Aug. 2023). <https://news.bloomberglaw.com/daily-labor-report/eeoc-settles-first-of-its-kind-ai-bias-lawsuit-for-365-000>
- [27] C. A. E. Goodhart. 1984. Problems of Monetary Management: The UK Experience. In *Monetary Theory and Practice*. Macmillan Education UK, London, 91–121. https://doi.org/10.1007/978-1-349-17295-5_4
- [28] Lara Groves, Jacob Metcalf, Alayna Kennedy, Briana Vecchione, and Andrew Strait. 2024. Auditing Work: Exploring the New York City algorithmic bias audit regime. *ACM Conference on Fairness, Accountability and Transparency* (2024). <https://doi.org/10.1145/3630106.3658959>
- [29] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. 2020. "It's a scavenger hunt": Usability of Websites' Opt-Out and Data Deletion Choices. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–12. <https://doi.org/10.1145/3313831.3376511>
- [30] Iina Ikonen, Francesca Sotgiu, Aylin Aydinli, and Peeter W. J. Verlegh. 2020. Consumer effects of front-of-package nutrition labeling: an interdisciplinary meta-analysis. *Journal of the Academy of Marketing Science* 48, 3 (May 2020), 360–383. <https://doi.org/10.1007/s11747-019-00663-9>
- [31] Indeed. 2023. *The Indeed Global AI Survey: Your Guide to the Future of Hiring*. Technical Report. Indeed. <https://www.indeed.com/lead/the-indeed-ai-report?hl=en#form>
- [32] John P. A. Ioannidis. 2019. What Have We (Not) Learnt from Millions of Scientific Papers with *P* Values? *The American Statistician* 73, sup1 (March 2019), 20–25. <https://doi.org/10.1080/00031305.2018.1447512>
- [33] Sonia K. Kang, Katherine A. DeCelles, Andrés Tilcsik, and Sora Jun. 2016. Whiteden Resúmes: Race and Self-Presentation in the Labor Market. *Administrative Science Quarterly* 61, 3 (Sept. 2016), 469–502. <https://doi.org/10.1177/0001839216639577> Publisher: SAGE Publications Inc.

- [34] Louis Kaplow and Steven Shavell. 1994. Optimal Law Enforcement with Self-Reporting of Behavior. *Journal of Political Economy* 102, 3 (June 1994), 583–606. <https://doi.org/10.1086/261947>
- [35] Kate Kaye. 2019. New York Just Set a ‘Dangerous Precedent’ on Algorithms, Experts Warn. *Bloomberg.com* (Dec. 2019). <https://www.bloomberg.com/news/articles/2019-12-12/nyc-sets-dangerous-precedent-on-algorithms>
- [36] Lauren Kirchner. 2017. New York City Moves to Create Accountability for Algorithms. *ProPublica* (Dec. 2017). <https://www.propublica.org/article/new-york-city-moves-to-create-accountability-for-algorithms>
- [37] Andreas Leibbrandt and John A. List. 2018. *Do equal employment opportunity statements backfire? Evidence from a natural field experiment on job-entry decisions*. Technical Report. National Bureau of Economic Research. <https://www.nber.org/papers/w25035>
- [38] M. Stuart Madden. 1986. The Duty to Warn in Products Liability: Contours and Criticism. *West Virginia Law Review* 89 (1986), 221. <https://heinonline.org/HOL/Page?handle=hein.journals/wvrb89&id=235&div=&collection=>
- [39] Scott de Marchi and James T. Hamilton. 2006. Assessing the Accuracy of Self-Reported Data: an Evaluation of the Toxics Release Inventory. *Journal of Risk and Uncertainty* 32, 1 (Jan. 2006), 57–76. <https://doi.org/10.1007/s10797-006-6666-3>
- [40] Arunesh Mathur, Gunes Acar, Michael J. Friedman, Eli Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. 2019. Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 1–32. <https://doi.org/10.1145/3359183>
- [41] Arunesh Mathur, Mihir Kshirsagar, and Jonathan Mayer. 2021. What Makes a Dark Pattern... Dark?: Design Attributes, Normative Considerations, and Measurement Methods. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. <https://doi.org/10.1145/3411764.3445610>
- [42] Alecia M. McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *Isjlp* 4 (2008), 543. https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/isjlpoc4§ion=27 Publisher: HeinOnline.
- [43] Jacob Metcalf. 2023. What federal agencies can learn from New York City’s AI hiring law. <https://thehill.com/opinion/technology/4360523-what-federal-agencies-can-learn-from-new-york-citys-ai-hiring-law/>
- [44] Jacob Metcalf, Ranjit Singh, Emanuel Moss, Emnet Tafesse, and Elizabeth Anne Watkins. 2023. Taking Algorithms to Courts: A Relational Approach to Algorithmic Accountability. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 1450–1462. <https://doi.org/10.1145/3593013.3594092>
- [45] Richard J. Murnane and John B. Willett. 2010. *Methods matter: Improving causal inference in educational and social science research*. Oxford University Press.
- [46] Arvind Narayanan. 2019. How to recognize AI snake oil. *Arthur Miller Lecture on Science and Ethics* (2019). <https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf> Publisher: Massachusetts Institute of Technology.
- [47] Don Norman. 2013. *The design of everyday things: Revised and expanded edition*. Basic books.
- [48] New York City Department of Consumer and Worker Protections. 2023. Automated Employment Decision Tools: Frequently Asked Questions. <https://www.nyc.gov/assets/dca/downloads/pdf/about/DCWP-AEDT-FAQ.pdf>
- [49] New York City Department of Consumer and Worker Protections. 2023. Automated Employment Decision Tools (Updated). <https://rules.cityofnewyork.us/rule/automated-employment-decision-tools-updated/>
- [50] Code of Federal Regulations. 2024. 29 CFR Part 1607 - General Principles. <https://www.ecfr.gov/current/title-29/part-1607/subject-group-ECFRdb347e844acdea6>
- [51] Robert N. Proctor and Londa Schiebinger. 2008. Agnotology: The making and unmaking of ignorance. (2008). <https://philarchive.org/archive/PROATM>
- [52] Mostafa Purmehdi, Renaud Legoux, François Carrillat, and Sylvain Senecal. 2017. The Effectiveness of Warning Labels for Consumers: A Meta-Analytic Investigation into Their Underlying Process and Contingencies. *Journal of Public Policy & Marketing* 36, 1 (April 2017), 36–53. <https://doi.org/10.1509/jppm.14.047> Publisher: SAGE Publications Inc.
- [53] 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. <https://doi.org/10.1145/3351095.3372828> arXiv:1906.09208 [cs].
- [54] Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. The fallacy of AI functionality. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 959–972. <https://dl.acm.org/doi/abs/10.1145/3531146.3533158>
- [55] J. Richard, Judith D. Singer, John B. Willett, and Richard J. Light. 2009. *By design: Planning research on higher education*. Harvard University Press.
- [56] David J. Rothman. 2017. *Strangers at the bedside: A history of how law and bioethics transformed medical decision making*. Routledge. <https://www.taylorfrancis.com/books/mono/10.4324/9781315130286/strangers-bedside-david-rothman>
- [57] Rebecca Schmidt and Colin Scott. 2021. Regulatory discretion: structuring power in the era of regulatory capitalism. *Legal Studies* 41, 3 (2021), 454–473.
- [58] Andrew D Selbst. 2021. An Institutional View of Algorithmic Impact Assessments. *Harv. J.L. & Tech.* 35 (2021), 78. <https://ssrn.com/abstract=3867634>
- [59] Jay P. Shimshack. 2014. The Economics of Environmental Monitoring and Enforcement. *Annual Review of Resource Economics* 6, 1 (Nov. 2014), 339–360. <https://doi.org/10.1146/annurev-resource-091912-151821>
- [60] Jay P. Shimshack and Michael B. Ward. 2008. Enforcement and over-compliance. *Journal of Environmental Economics and Management* 55, 1 (Jan. 2008), 90–105. <https://doi.org/10.1016/j.jeem.2007.05.003>
- [61] Jodi L. Short and Michael W. Toffel. 2016. The integrity of private third-party compliance monitoring. *Admin. & Reg. L. News* 42 (2016), 22. https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/admreln42§ion=11 Publisher: HeinOnline.
- [62] SHRM. 2022. Automation & AI in HR. https://advocacy.shrm.org/SHRM-2022-Automation-AI-Research.pdf?_ga=2.112869508.1029738808.1666019592-61357574.1655121608
- [63] Michelle Singletary. 2023. The FTC wants to make it easy to cancel subscriptions - The Washington Post. *The Washington Post* (March 2023). <https://www.washingtonpost.com/business/2023/03/24/ftc-cancel-subscriptions/>
- [64] Jacob Soll. 2014. *The reckoning: Financial accountability and the rise and fall of nations*. Basic Books (AZ).
- [65] Daniel J. Solove. 2012. Introduction: Privacy self-management and the consent dilemma. *Harv. L. Rev.* 126 (2012), 1880. https://heinonline.org/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/hlr126§ion=87 Publisher: HeinOnline.
- [66] Daniel J. Solove. 2021. The Myth of the Privacy Paradox. *George Washington Law Review* 89 (2021), 1. <https://heinonline.org/HOL/Page?handle=hein.journals/gwlr89&id=15&div=&collection=>
- [67] Richard H Thaler and Cass R Sunstein. 2003. Libertarian Paternalism. *American Economic Review* 93, 2 (April 2003), 175–179. <https://doi.org/10.1257/000282803321947001>
- [68] Peter V. B. Unger, Karen Van Esen, Alexander Birkhold, Mohammed T. Farooqui, and Elizabeth Satarov. 2023. DOJ Voluntary Self-Disclosures Safe Harbor Policy. (Oct. 2023). <https://natlawreview.com/article/doj-announces-new-mergers-acquisitions-safe-harbor-policy-voluntary-self>
- [69] 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. <https://doi.org/10.48550/arXiv.2202.09519> arXiv:2202.09519 [cs].
- [70] Christine Webber. 2023. AI Bias Panel Shows EEOC Should Ditch Four-Fifths Rule – Law360 Expert Analysis | Cohen Milstein. <https://www.cohenmilstein.com/update/ai-bias-panel-shows-eeoc-should-ditch-four-fifths-rule-%E2%80%93-law360-expert-analysis>
- [71] John T. Wilbanks. 2020. Electronic Informed Consent in Mobile Applications Research. *The Journal of Law, Medicine & Ethics: A Journal of the American Society of Law, Medicine & Ethics* 48, 1_suppl (March 2020), 147–153. <https://doi.org/10.1177/1073110520917040>
- [72] Meg Young, Michael Katell, and P.M. Krafft. 2022. Confronting Power and Corporate Capture at the FAccT Conference. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. ACM, Seoul Republic of Korea, 1375–1386. <https://doi.org/10.1145/3531146.3533194>
- [73] Jonathan Zong. 2020. From individual consent to collective refusal: changing attitudes toward (mis)use of personal data. *XRDS: Crossroads, The ACM Magazine for Students* 27, 2 (Dec. 2020), 26–29. <https://doi.org/10.1145/3433140>
- [74] Jonathan Zong and J. Nathan Matias. 2022. Bartleby: Procedural and Substantive Ethics in the Design of Research Ethics Systems. *Social Media + Society* 8, 1 (Jan. 2022), 205630512210770. <https://doi.org/10.1177/20563051221077021>
- [75] Jonathan Zong and J. Nathan Matias. 2023. Data Refusal From Below: A Framework for Understanding, Evaluating, and Envisioning Refusal as Design. *ACM Journal on Responsible Computing* (Oct. 2023), 3630107. <https://doi.org/10.1145/3630107>

ETHICS AND ADVERSE IMPACT STATEMENT

Ethical Review

We submitted this study for ethics review to Cornell University in protocol #IRB0147978. Because the subjects in this study are organizations and the data is public, they concluded that the study does not qualify as human subjects research regulated under the Common Rule. We have nonetheless taken measures to minimize risk, contacting employers to give them a chance to correct any errors. We have also de-coupled the identities of student investigators from the dataset.

Adverse Impacts

It is possible that this paper, when published, could influence and inform wider policy conversations alongside news coverage and enforcement of Local Law 144. Since this study focuses on organizations, we adopted best practices from journalism [23] to ensure the accuracy and reliability of our data, giving employers the opportunity to correct any errors.

A APPENDIX

A.1 Key Terms for understanding employment discrimination

Selection Rate: The rate at which members of a defined demographic group advance in a hiring process. "Selection" is any decision point at which candidates are either chosen to move to the next stage—from resume screening to hiring—or rejected; it does not refer only to the final hiring decision.

$$\frac{\text{Number of applicants in a category}}{\text{Total number of applicants}}$$

Scoring Rate: Scoring rate is the same concept as selection rate, but captures an important aspect of AEDTs generally: many AEDTs output a simplified *scoring* or *ranking* rather than a *decision*. Examples include personality scores, culture fit scores, intelligence rankings, etc. The formula in LL 144 accounts for the rate at which the system scores applicants from the protected group with *above median*, or desirable, scores.

$$\frac{\text{Number of applicants in a category with score > median score}}{\text{Total number of applicants in category}}$$

Impact Ratio: AKA impact factor, is a measure of the relative selections rates between one group and the most selected group. It measures differences in rates, not in absolute numbers. Identical selection rates result in an impact ratio of 1.0. The lower the number, the more discriminatory the selection process is. An impact ratio may also be calculated comparing the rates for the less-selected group against the rates for the entire population of candidates.

$$\frac{\text{Selection rate for a category}}{\text{Selection rate of the most selected category}} \quad \text{OR} \quad \frac{\text{Scoring rate for a category}}{\text{Scoring rate of the highest scoring category}}$$

Disparate Impact: AKA adverse impact, is an impact ratio that is low enough to be illegal or otherwise impermissible. In theory, disparate impact could be set at any possible impact ratio. In practice, it is largely defined by/identical to the four-fifths rule. Disparate impact is used to capture the phenomenon of unintentional discrimination through systemic features. An employer may be legally responsible for causing disparate impact even if none of their procedures explicitly discriminate against a protected group, e.g., if an employer has no policy against hiring a certain gender but their hiring practices result in *relatively* lower rates of hiring for that gender, then they *may* be guilty of causing disparate impact. Disparate impact may measure unintentional discrimination, but it does *not* measure historical discrimination. The numerator may—and usually does in practice—correspond to a historically disadvantaged ‘protected category,’ but it is not synonymous. Impact ratios can measure discriminatory outcomes against non-protected/historically-advantaged group. For example, an employer may have a policy of only hiring members of a historically disadvantaged group for a low-status job, resulting in a low disparate impact ratio when the historically disadvantaged group is the numerator (which would likely be illegal discrimination against both groups).

Four-Fifths Rule: AKA 0.8 Rule or 80% Rule, is a convention in US anti-discrimination law that defines disparate impact as any impact ratio that falls below 0.8, or four-fifths. In other words, if the rate at which a group is selected is lower than 80% percent of the rate at which the most-selected group is selected, then the impact is presumed ‘disparate’ or ‘adverse.’ The ‘rule’ in the name is a misnomer, it is instead a guideline that grants some degree of protection against regulatory scrutiny and litigation for hiring practices that result in impact ratios above 0.8. Falling below an impact ratio of 0.8 is not considered to be automatic evidence of illegal discrimination absent an investigation that considers many other factors that may legally justify a discriminatory outcome, such as legitimate relevance of the features to the job. However, falling below an impact ratio of 0.8 is a strong signal that a hiring process *could* receive scrutiny and requires additional justification. The EEOC defines the Four-Fifths Rule in CFR §1607.4(D)¹⁵.

$$\frac{\text{Selection rate for a category}}{\text{Selection rate of the most selected category}} \leq 0.8$$

A.2 Key Legal and Historical Details

Despite being a municipal law, the relationship between LL 144 and federal anti-discrimination law is critically important to the outcomes of this study. The so-called “four-fifths rule” [3, 50, 69] that emerged from US anti-discrimination law looms large in employment jurisprudence and algorithmic fairness metrics. The “four-fifths rule” is a general guideline that an impact ratio greater than 0.8 is considered *de facto* evidence (but not *de jure* evidence) of non-discriminatory outcomes. The concept of disparate impact and the four-fifths rule loom large in discussions of algorithmic bias because the unintentional, systemic discrimination described by disparate impact is often considered structurally analogous to algorithmic discrimination [2]. Indeed, in many algorithmic bias detection and mitigation tools, the four-fifths rule is treated as a foundational, standard measure of the line between acceptable and unacceptable discrimination, despite the fact that it is domain-specific inside of US law. As critics have noted [69], the four-fifths rule is a fundamentally arbitrary threshold not supported by robust empirical evidence. But it is historically important, well-litigated, and embedded in U.S. anti-discrimination jurisprudence, and therefore familiar and unavoidable.

LL 144 makes no explicit mention of a 0.8 threshold or any other threshold. As the DCWP states in a FAQ, “the law does not require any specific actions based on the results of a bias audit, including ceasing the use of an AEDT shown to result in disparate impact” [49]. The FAQ does note that other laws on employment discrimination still apply outside of the scope of LL 144, which would presumably include the EEOC’s definition of disparate impact [14, 48, 50].¹⁶ Because LL 144 does not take a stance on permissible thresholds, it is also silent on remediation, and thereby limits the utility of the law for driving and measuring reduction of discrimination over time.

¹⁵Available here: <https://www.ecfr.gov/current/title-29/subtitle-B/chapter-XIV/part-1607/subject-group-ECFRdb347e844acde6>

¹⁶The first case of algorithmic hiring bias—for age discrimination—was settled with the EEOC in 2023, but it was for a system that automatically screened out all older applicants and not a typical disparate impact scenario [15, 26].

Federal regulators have begun to set expectations about algorithmic systems and discrimination law that forbid employers from creating adverse impact through AI systems [11, 13, 14], including holding hearings on the appropriateness of the four-fifths rule for AI use cases [16, 70]. Nonetheless, this guidance is situated in the EEOC’s long-standing decision tree for determining acceptable selection practices, which offers few hard lines amenable to auditing and relies heavily on judgment calls and precedents [12]. Nor has any federal body established guidance about safe harbors for audits conducted in compliance with a local or state law. If introduced, such safe harbors would protect employers from disincentives to disclose information about discriminatory algorithms, as is the case in other areas of law [68]. Given the absence of any safe harbors, employers complying with LL 144 by posting audit reports might also be opening themselves to other types of liability.

A.3 Survey Response Rates

Of all the employers in our sample, 29 listed no email or contact form or had a contact form with a character limit that was so low as to make the message incomplete. Of the emails we did send in the first round, 15 bounced or received a reply saying the address was not monitored.

We gave the employers 22 days to respond to our second outreach, from Dec 19 - Jan 10th. Among the 387 follow-up survey emails that we sent, 26 employers completed the survey, and 15 contacted us in some other way, either by email or phone. In total, this follow-up process yielded two audit reports that we had not previously observed.

A.4 Student Investigators

We are deeply grateful to the students of COMM/INFO 2450 for their thorough and extensive collaboration on this project, among them Amelia Neumann, Andrew Wu, Angelina Chen, Anjiya Amrani, Anushka Shorewala, Bella Samtani, Bingsong Li, Carina Wang, Carolina Michailoff, Chelsea Lin, Chengling Zheng, Diana Flores Valdivia, Doan-Viet Nguyen, Dora Xu, Erik Starling, Evelyn C. Kim, Gianna Chan, Haley Qin, Hannah M. Yeh, Hermione Bossolina, Hope Best, Ingrid Gruener Luft, Jacob Levin, Jimin Kim, Jolene Ie, Kashmala Arif, Katherine Hahnenberg, Kathryn M. Papagianopoulos, Kevin Jianzhi Wang, Kexin Li, Kimmie Jimenez, Lili Mkrtychyan, Lindsay Peck, Maksym "Max" Bohun, Mark Timothy Bell, Mika Labadan, Minh H. Le, Neha Sunkara, Nicholas Bergersen, Nicholas Won, Nicole Tian, Noah Salzman, Nuo Cen, Omar Ahmed, Owen J. Chen, Reid Fleishman, Reinesse Wong, Sebastian Klein, Shukria Mirzaie, Simah Sahnosh, Siying Cui, Sophia Torres Lugo, Sritanay Vedartham, Subhadra Das, Thej Khanna, Varsha Gande, Weiyan Zhang, Wen Yu Chen, Yanran Li, Yiwen Zhang, Yubin Yang, Yuchen Yang, Yuyan Wu, and Zoey Arnold.