Auditing Effects on Employment Hiring: Evidence from the New York City Algorithmic Bias Audit Law[1]

Daniel Aobdia (daniel_aobdia@psu.edu)
The Pennsylvania State University

Hao Ma (hao.ma@hhs.se)
Stockholm School of Economics

Sheryl Zhang (zhangs@essec.edu)
ESSEC Business School

This draft: May 2026

Abstract

This paper examines the effect of algorithmic bias audits on workforce diversity by exploiting New York City's (NYC) pioneering Bias Audit Law (Local Law 144). The law requires third-party audits for NYC employers that use algorithm-driven hiring tools to examine instances of disparate impact in recruiting. Using a difference-in-differences analysis, we find that bias audit requirements lead to almost a two percentage-point reduction in the share of male hiring, while having no significant effects on white employee hiring. The effects of decreased male hires are particularly pronounced among firms that directly or indirectly obtain an audit and have greater pre-audit gender imbalances. The decline is also concentrated in lower to mid-tier positions and applies to both public and private companies. We also document costs in the form of increased human resource staff hires and longer times to fill vacancies. Finally, we find evidence of subsequent improvement in performance using financial analysts' forecast accuracy as a setting. Our results provide the first large-scale empirical evidence on the effectiveness of mandated algorithmic bias audits, highlighting a trade-off between levelling the playing field for job seekers and reduced recruitment efficiency.

Keywords: Algorithmic bias; Artificial Intelligence; Bias audits; Workforce diversity
JEL Classification Codes: J71; L51; M41; M51; O33

---

[1] Daniel Aobdia was a Senior Economic Research Fellow in the Center of Economic Analysis at the PCAOB between September 2014 and September 2016. He currently advises the PCAOB on questions related to the economics of auditing. The views expressed in this paper are the views of the authors and do not necessarily reflect the views of the board, individual board members, or staff of the PCAOB. We appreciate insightful feedback from Anthony Le, Bin Ke (discussant), Rick Laux, Vincent Lin, Henock Louis, Kelvin Law, Xiumin Martin, Robert Stoumbos, and workshop participants at the 6th Annual Labor and Accounting Group Conference, the 2026 Nanyang Business School Accounting Conference, Bayes Business School, Oklahoma State University Distinguished Speaker Lecture Series, Shanghai Jiaotong University Antai Business College, and the Stockholm School of Economics. All errors are our own.

"Automated employment decision tools are in urgent need of transparency, oversight, and regulation. These technologies all too often replicate and amplify bias, discrimination, and harm towards marginalized communities."

Daniel Schwarz, senior privacy and technology strategist at the New York Civil Liberties Union

1. Introduction

In 2024, roughly 80% of U.S. companies and nearly all Fortune 500 companies employed automated systems to screen, score, and rank job market applicants (Dennison 2024). While Artificial Intelligence (AI)-enabled recruitment has the potential to enhance recruitment quality, increase efficiency, and reduce transactional work, it could also pose a significant threat by denying equal access to jobs and discriminating against certain categories of applicants (e.g., Chen 2023; An et al. 2025). Such bias often arises from training data that reflects historical inequalities or underrepresentation, as well as from algorithmic design choices that reinforce discriminatory patterns (Barocas and Selbst 2016; Celiktutan et al. 2024). High-profile examples, especially Amazon's 2014 AI recruiting experiment that led to unintentional bias against women, highlight ongoing concerns about algorithmic bias in hiring processes (e.g., Dastin 2018; Goodman 2018).[2,3] In response to these concerns, New York City (NYC) implemented the Bias Audit Law on July 5, 2023.[4] This pioneering legislation requires employers using Automated Employment Decision Tools (AEDTs) to conduct independent annual bias audits and publicly disclose both the use of the AEDTs and the audit results. As a result, bias audit, systematic independent assurance of internal hiring algorithms for instances of selection disparities, has emerged as a leading policy mechanism to promote fairness and transparency in algorithmic

[2] Amazon ultimately cancelled the project. According to Dastin (2018), the bias existed "because Amazon's computer models were trained to vet applicants by observing patterns in resumes submitted to the company over a 10-year period. Most came from men, a reflection of male dominance across the tech industry… In effect, Amazon's system taught itself that male candidates were preferable."

[3] Other high-profile examples exist, such as Derek Mobley's collective action lawsuit alleging that Workday's AI screening disproportionately excluded candidates based on race and age (Gassam Asare 2025).

[4] We refer to these audits as bias audits in the remainder of the paper, based on their given names. However, we note that these audits typically aim to measure disparate impact, not necessarily bias. See section 2 for more details.

hiring. However, how bias audits affect algorithmic fairness in real-world hiring decisions and the ultimate composition of new hires and the workforce remain open questions. Practitioners, scholars, and policy experts hold sharply divided opinions on the effectiveness of bias audits, underscoring the urgency of evaluating their effectiveness (e.g., Kroll et al. 2017; Leicht-Deobald et al. 2019; Moschella 2022; Celiktutan et al. 2024; Hilliard et al. 2024; Lam et al. 2024). This paper addresses this critical question by examining the effect of bias audit requirements on hiring diversity.

On the one hand, proponents of bias audits claim that auditing for discrimination is an essential part of the strategy for detecting and responding to biased algorithms. They draw a direct parallel to financial-statement audits, where decades of empirical research show that third-party assurance substantially reduces misstatements (BABL AI 2022; Lam et al. 2024), enhances information quality, and aligns managerial incentives with stakeholder interests (Francis 2004; DeFond and Zhang 2014). By the same logic, systematic audits of hiring algorithms can help diagnose hidden group-based selection gaps. As Kim (2017) asserts, auditing increases firms' awareness of their algorithms' impact on hiring decisions, encouraging corrective actions and deterring firms from engaging in discriminatory hiring practices.

On the other hand, critics argue that bias audits only focus on disparate impact outcomes, and thus often treat the decision-making process as a black box, without identifying the root causes of disparities or assessing their substantive significance (Kroll et al. 2017; Leicht-Deobald et al. 2019). This flaw leaves firms without clear remediation guidance and opens the door to dismissing genuine fairness issues or applying superficial tweaks that leave underlying biases intact or even create new ones. Further, employers might simply scale back AI screening in favor of human decision-makers, potentially reintroducing unconscious stereotypes that might be more difficult to detect or remediate at scale (Moschella 2022; Celiktutan et al. 2024).

Despite this debate, there is little empirical evidence on the effect of bias audits. To fill this gap, we provide the first large-sample empirical evidence on the effectiveness of bias audits by investigating NYC's Bias Audit Law (Local Law 144), the first law in the U.S to require an AI bias audit. Effective on July 5, 2023, the law prohibits positions located in NYC from using an AEDT in hiring unless the tool has been subject to a public bias audit report within one year of the use of the tool, effectively requiring annual audits for tools in continuous use. These audits must be performed by third parties and check for instances of disparate impact built into these systems. At minimum, an independent auditor's evaluation must include calculations of selection or scoring rates and the impact ratio across sex, race/ethnicity, and intersectional categories.

We employ a difference-in-differences analysis to examine the effects of bias audits on the gender and racial composition of new hires. We are particularly interested in gender because the 2014 Amazon AI recruiting experiment resulted in unintentional bias against women, in part because Amazon operates in a male dominated profession (Dastin 2018). Using data from Revelio Labs, we construct office-level measures at the year-quarter level. Workforce diversity is measured by the proportion of male hires and white hires.[5,6] Offices in NYC constitute the treated group, while offices outside NYC, belonging to the same firms that have at least one office in the city, serve as controls. We then compare how hiring outcomes change differently between treated and control offices around the passage of the Bias Audit Law.

[5] While the Bias Audit Law specifically targets NYC, we use the New York MSA restricted to the New York State as our treated area due to data limitations. This approximation is reasonable given that: (1) NYC comprises approximately 70% of the New York State portion of the NY MSA by population and contains the majority of commercial districts; (2) the MSA represents a unified labor market with extensive cross-boundary commuting; and (3) most large employers subject to audit requirements maintain primary offices within NYC proper.

[6] One caveat of using the Revelio Labs data is that these data are mostly restricted to professional and white-collar positions. However, using these data also mitigate alternative explanations, such as a concurrent push for more union organizing.

We find that the adoption of bias audits leads to an almost two percentage-point drop in the share of male hires, representing a 3.2% reduction of the mean share. We also estimate the dynamic effect of the Bias Audit Law and document no systematic evidence of pre-trends, which provides confidence vis-à-vis the parallel trend assumption. An intersectional breakdown shows that this reduction in male hiring occurs uniformly across racial lines: both white-male and non-white-male cohorts experience statistical decreases in their representation. This finding underscores the law's intent to address overlapping dimensions of bias rather than targeting a single demographic group. In addition, we do not find evidence to suggest that the aggregate share of white hires changes. One possible explanation is that both the Revelio Labs dataset and resume screenings infer race from individuals' names, and name-based algorithms are generally less accurate in predicting race than gender (Lockhart et al., 2023). Evidence from the Amazon case also indicates that bias against women in AI-based hiring algorithms might be common, especially in male dominated professions.[7]

To ensure that these results isolate the plausibly causal impact of the Bias Audit Law, and are not driven by NYC's unique culture and initiatives or by other labor□market shifts, we perform a placebo test on age cohorts: since the law explicitly targets gender and racial categories, it should have no effect on the age composition of hires, which is not covered by the law. As expected, we observe no significant changes in hiring shares across different age groups, confirming that the gender□race findings are not confounded by contemporaneous policies such

[7] Subsequent granular analysis of the non-white subpopulation in the online appendix reveals statistically significant increases in Black hiring alongside statistically significant declines in Hispanic hiring. These offsetting movements account for the net-zero change in the broad white-hire metric, while still reflecting substantive redistribution of opportunities among non-white groups. Together, these results demonstrate that bias audits can effectuate broad reductions in overrepresentation for some groups and nuanced shifts in racial composition, even when aggregate measures appear unchanged.

as pay-transparency mandates or post-pandemic labor-supply shifts.

There are two possible mechanisms that drive our results. First, bias audits can lead to identification and remediation of issues that result in disparate impact. For this mechanism to apply, a firm would need to explicitly undergo a bias audit. Second, companies could scale back AI recruiting initiatives and increase the level of human oversight, with an explicit focus on reducing potential biases. Doing so could result in some firms being exempt from bias audits. We conduct a series of analyses to explore these two mechanisms in more detail and find evidence consistent with both mechanisms explaining our main result.

First, we identify a set of companies that directly or indirectly undergo bias audits. We find that while a limited number of firms directly commission bias audits, similar to the findings of Wright et al. (2024), the incidence of indirect audits, for which leading AI vendors themselves conduct bias audits that ultimately impact their customers, is much higher.[8] We estimate that 39% of the NYC firms in our sample either directly receive a bias audit or receive one indirectly because their AI vendor received a bias audit.[9] Importantly, our results are stronger and often concentrate in offices that directly or indirectly obtain a bias audit. The effect is particularly pronounced when we focus on algorithms that exhibited higher mean male hiring rates prior to the law's enactment. This finding is consistent with the idea that algorithms trained on more gender-imbalanced populations exhibit greater bias, a pattern also reflected in the anecdotal evidence from Amazon (Dastin 2018; Goodman 2018).

Second, we find stronger results, concentrated among offices that obtain a bias audit, when

[8] We find that many companies use off the shelf AI recruiting solutions from data vendors such as Workday. Thus, they can benefit from their data vendors conducting a bias audit. These data vendor audits are similar in spirit to service organization control (SOC) audits received by service organizations on behalf of their customers (e.g., AICPA 2023; Schoenfeld 2024).

[9] Incidences of firms indirectly obtaining bias audits are much greater than firms directly obtaining such audits. Out of the 2,522 companies in our primary sample, we estimate that 15 (975) directly (indirectly) obtain an audit.

these offices have limited human resources (HR) support. This result is consistent with the use of AI algorithms resulting in greater bias when human oversight is limited, and thus there is greater reliance on AI systems (e.g., Skitka et al. 1999; Wilson et al. 2025). To further understand whether greater HR and thus human oversight can help reduce bias, we explore and find an increase in the net hiring rate of HR staff, both for companies that obtain and do not obtain an audit. This result is consistent with companies resorting to more manual hiring processes following the Bias Audit Law. In addition, our results on male hire reduction and increased HR hiring concentrate in male dominated industries and firms, where potential AI bias was greater prior to the Bias Audit Law. These results apply for both companies that obtain and do not obtain an AI bias audit. Overall, these results are consistent with firms, especially those for which gender imbalances are greater and thus where bias might be more prevalent, resorting to greater human oversight in reaction to the law. In some instances, doing so allows these firms to avoid having to undergo an AI Bias Audit.

We conduct several additional analyses that provide additional insights. First, we examine whether the effects are limited to public firms, which face greater reputational and regulatory scrutiny. We find that both public and private firms reduce male hiring following the law's implementation. Second, we test whether the effects vary by career level. The decline in male hiring is particularly pronounced in lower to mid-tier positions, likely because these roles rely more on AI-based employment tools (Godfrey 2025; Jaser et al. 2021) and involve lower adjustment costs in response to fairness mandates (Chang and Kirgios 2024; Larcker et al., 2025). Third, we explore potential spillover effects for firms that directly or indirectly obtain AI audits and find evidence of spillover effects for non-NYC offices of such firms. These results are consistent with firms using similar hiring processes across different offices. However, the male hire reduction is less pronounced outside of NYC. This is consistent with firms not always following the same hiring processes across different offices, or taking additional actions in their NYC offices, such as increasing HR staff, that ultimately impact the proportion of male hires.

We also examine additional costs of implementing the Bias Audit Law besides an increase in HR staff. Critics argue that the law may delay hiring processes, prolong staffing shortages, and disrupt overall recruitment operations (Wade 2022). Indeed, mandating third-party audits could create both administrative overhead (e.g., commissioning reports and coordinating demographic data) and pressure to extend searches to underrepresented groups to improve diversity performance, potentially reducing recruitment efficiency (Matsa and Miller 2013). Consistent with these concerns, we document a significant increase in the average days to fill a vacancy among treated offices, capturing potential inefficiencies introduced by compliance requirements (Chen and Li 2023). However, we observe no change in the total number of hires, suggesting that firms absorb the delay without curtailing overall recruitment volume and shifting their workforce to offices not subject to the Bias Audit Law, including offices in the New York MSA but located in states that do not have bias audit requirements.

So far, our results show changes in the composition of new hires, but they might also be explained by companies making arbitrary changes in response to the law that aim to reduce disparate impact but not necessarily reduce actual bias in recruiting.[10] To disentangle between both possibilities, we need to investigate whether changes in hiring outcomes following the Bias Audit Law translate into broader firm performance. Improvements in firm performance would be consistent with the law leading to the hiring of more qualified individuals that ultimately result in

[10] For example, disparate impact could be explained by idiosyncratic differences in average applicant qualifications in specific categories. A greater selection rate of the more qualified applicants might thus not reflect greater bias, but disparate impact could be present because of the average skill imbalance in the pool of candidates.

improved firm outcomes. This question, however, is difficult to answer because individual office-level performance is not available for most companies. To shed more light on this, we first focus on analyst forecast accuracy, using a sample of brokerage firms located in and outside of NYC. We confirm similar employment patterns for brokerage firms: NYC offices reduce the share of male hires in response to the Bias Audit Law, increase HR staff hiring, and increase the number of days to fill vacancies. Importantly, we also find that analysts in NYC offices exhibit greater forecast accuracy, a proxy for performance, during the post period (e.g., Clement 1999; Mikhail et al. 1999; Loh and Mian 2006). Again, these results are concentrated among offices that obtain bias audits. This evidence thus suggests that brokerage firms hire more talented individuals who help analysts provide more accurate forecasts, consistent with potential biases existing prior to the implementation of the law.

Second, we use new hire perceptions of their workplace, measured by employee satisfaction ratings. We find increased employee satisfaction in treated offices for overall ratings, diversity and inclusion ratings, and career opportunity ratings. Again, these results concentrate among offices that obtain bias audits. Collectively, these results provide initial evidence that beyond reducing disparate impact, the law may also help reduce underlying biases.

This paper makes several contributions. First, by providing the first large-scale, plausibly causal evidence on the effects of bias audits, our study offers guidance for policymakers considering similar mandates. We find that third-party audits help rebalance hiring while introducing additional costs in recruitment timelines. These findings can inform the calibration of proposed bias-audit legislation, such as New Jersey's Senate Bill S1588 and Pennsylvania's House Bill HB 594. Beyond immediate policy applications, this research advances the fair AI governance debate by illuminating the practical consequences and tradeoffs that arise when regulating algorithmic hiring technologies to level the playing field and create more equality of chances among different populations.

Second, this study contributes to the emerging literature on AI in corporate decision-making by providing the first evidence of regulatory interventions targeting algorithmic bias in hiring. While recent research has examined AI adoption across various corporate functions, including financial reporting, financial auditing, investment decisions, and capital markets (Bradshaw et al. 2025; Chang et al. 2025; Blankespoor et al. 2025; Choi and Xie 2025; Levy 2025; Munoko et al. 2020), the intersection of AI governance and human capital management remains largely unexplored. Our work fills this gap by investigating how regulatory oversight can constrain algorithmic bias in hiring, extending the study of AI governance from financial markets to labor markets and contributing to a more comprehensive understanding of how firms can manage AI-related risks across operational domains.

Third, this study extends the scope of third-party auditing beyond financial oversight to employment practices, contributing to an emerging literature on the expanding role of auditors in corporate governance. Building on foundational work on audit independence and quality (Watts and Zimmerman 1983; DeAngelo 1981; DeFond and Zhang 2014; Lawrence et al. 2011), we demonstrate that the monitoring and verification functions central to financial auditing can be effectively applied to algorithmic hiring systems. Our findings complement research on audit scope expansion into non-financial domains, such as sustainability assurance (Aobdia et al. 2025; Gipper et al. 2025) and DeFi assurance (Bourveau et al. 2024; Knechel, Maex, and Park 2025), by showing that employment bias audits generate meaningful behavioral changes in organizational hiring practices.

Finally, this study contributes to the literature on gender diversity. Prior studies have documented the consequences of gender diversity on firms' financial outcomes (e.g., Lins et al., 2024; Liu et al., 2023; Breuer et al., 2024; Billings et al., 2022; Adhikari et al., 2019). Scholars have also explored the mechanisms that promote or hinder progress in corporate diversity, such as same-group representation (Dong 2022; Ahn et al., 2025), diversity quotas or targets (Cai et al., 2024), occupational licensing requirements (Sutherland et al. 2024), disclosure mandate (Bourveau et al., 2025), and diversity information in job postings (Choi et al., 2023). While this literature underscores the importance of institutional design and regulatory levers, much less is known about how emerging technologies, particularly AI-driven hiring tools, affect gender diversity, and whether existing governance mechanisms are sufficient to counteract potential biases. Our study complements this literature by showing that auditing can serve as a mechanism that mitigates algorithm-driven gender bias in hiring, adding to the organizational and regulatory tools that promote fairness in hiring.

2. Institutional Background of New York City's Local Law 144

NYC's Local Law 144 represents the first comprehensive municipal regulation of AI bias in employment decisions in the United States. Effective July 5, 2023, the law prohibits employers and employment agencies from using AEDTs for positions located in NYC, unless these tools have undergone independent bias audits within the previous year, effectively mandating annual audits for any tools in continuous use. Importantly, the law does not require any specific remedial actions based on audit results, focusing instead on increasing transparency. Figure 1 presents the timeline of the NYC Bias Audit Law.

[Insert Figure 1 about here]

An AEDT is defined as a computer-based tool that uses machine learning, statistical modeling, data analytics, or AI to help employers make employment decisions and substantially assist or replace discretionary decision-making. This broad definition encompasses various AI and algorithmic systems used in recruitment, screening, and promotion processes.

The law establishes three core requirements: mandatory independent bias audits conducted annually by third parties, notification to candidates about AEDT usage at least 10 business days prior to evaluation, and public disclosure of audit results on company websites for at least six months. At minimum, bias audits must examine selection rates and impact ratios for gender/sex (male, female, and optionally other) and race/ethnicity categories (Hispanic or Latino, White, Black or African American, Native Hawaiian or Pacific Islander, Asian, Native American), with specific attention to intersectional bias across these dimensions. The selection rates and impact ratios are calculated as follows:

$$selection\ rate = \frac{number\ of\ applicants\ selected\ by\ AEDT\ from\ a\ category}{number\ of\ total\ applicants\ from\ a\ category},$$

and

$$Impact\ ratio = \frac{selection\ rate\ for\ a\ category}{selection\ rate\ of\ the\ most\ selected\ category(eg.Male, white\ \ldots.)}.$$

Impact ratios below 0.80 are considered evidence of potential adverse impact under the federal "4/5ths rule," though the law *does not* mandate specific remedial actions when this threshold is crossed. In practice, bias audits are conducted by various types of third parties, including specialized consulting firms (such as Ambrose Consulting, LLC), statistical and economic expert consulting services (such as BLDS, LLC), and algorithmic auditing firms (such as BABL AI Inc.). These auditors typically possess expertise in employment discrimination analysis, statistical modeling, and algorithmic fairness assessment. Bias audits follow a standardized methodology that includes comprehensive system descriptions, methodological frameworks specifying data collection periods and sample sizes, demographic analysis across required categories, statistical calculations of selection and impact ratios, and organizational governance assessments. Online Appendix A provides a sample audit report.

Enforcement is handled by the NYC Department of Consumer and Worker Protection (DCWP), with civil penalties ranging from \$500 to \$1,500 per violation and private rights of action available to affected candidates. The law applies to employers using AEDT for positions physically based in the city, regardless of where employers are headquartered. The DCWP clarifies that the key determinant is where the position is physically "located." Bias audits must be conducted for positions located within NYC, whereas roles based elsewhere are exempt. Fully remote positions trigger audit requirements only when the employer's sole offices reside in the city; employers with no NYC offices face no audit obligation, and those operating both inside and outside the city must determine applicability by fact-specific analysis.

Several features make Local Law 144 particularly well-suited for causal identification. The law's geographic specificity creates clear treatment boundaries between NYC and other locations, while its mandatory nature and specific effective date provide clean variation for difference-in-differences analyses. Organizations with operations both inside and outside NYC face identical management and corporate policies, differing only in their exposure to the bias audit requirement.

3. Literature Review and Hypothesis Development

3.1 AI in Corporate Decision-Making and Hiring

The integration of AI in corporate decision-making has fundamentally transformed how organizations operate across multiple domains. Recent literature documents AI's expanding role in financial reporting (Bradshaw et al. 2025), investment decisions (Choi and Xie 2025), and capital markets (Levy 2025), establishing a foundation for understanding AI's broader organizational implications. Within human capital management, this technological shift has been particularly pronounced, with approximately 80% of U.S. companies employing automated systems to screen, score, and rank job applicants in 2024 (Dennison 2024).

There are three key advantages to using AI-driven recruiting. First, sophisticated pattern-recognition algorithms can enhance selection quality by identifying subtle predictors of job performance while minimizing the influence of irrelevant factors and personal biases that affect human screeners (McColl and Michelotti 2019). Second, automating time-intensive processes such as resume parsing and initial screening significantly improves operational efficiency, reducing manual review time and enabling recruiters to focus on higher-value activities (Van den Broek et al. 2021). Third, algorithmic systems can help standardize hiring procedures by applying consistent, objective criteria across all candidates, theoretically reducing the impact of individual prejudices and increasing fairness in recruitment decisions (Raghavan et al. 2020).

However, this optimistic view confronts increasing evidence of algorithmic bias in real-world hiring. For example, An et al. (2025) instruct commonly used large language models, such as ChatGPT, Claude and Gemini, to score 361,000 resumes and find significant biases that may shift hiring probabilities by one to three percentage points. In a salient instance, Amazon discontinued a 2014 project to automate hiring because the tool inadvertently became biased against women (Dastin 2018; Goodman 2018). High-profile legal challenges, including Derek Mobley's collective action against Workday alleging that AI screening systems disproportionately excluded candidates based on race and age, also underscore the gap between theoretical potential and practical implementation.

To understand why these high-profile failures occur, the literature identifies two primary sources of discrimination in AI hiring systems. The first source stems from data quality problems: algorithms produce biased outcomes when trained on inaccurate (Kim 2016), biased (Barocas and Selbst 2016), or unrepresentative input data (Suresh and Guttag 2021). This creates a problematic feedback loop where historical hiring discriminations become systematically embedded in algorithmic decisions, perpetuating past inequities rather than correcting them. This mechanism likely drove the Amazon project issues (Dastin 2018). The second mechanism involves proxy discrimination, where seemingly neutral variables, such as zip codes, educational institutions, or word choices in resumes, serve as statistical proxies for protected characteristics like race or gender (Prince and Schwarcz 2019). Even when algorithms explicitly exclude protected attributes, these correlated variables enable indirect discrimination that may be difficult to detect or prove in legal contexts. Taken together, these mechanisms suggest that without active intervention, AI systems may perpetuate employment discrimination.

3.2 Policy Interventions in Labor Markets and Employment Discrimination

The economics literature on employment discrimination provides crucial context for understanding regulatory responses to algorithmic bias. Building on foundational work by Becker (1957) and Arrow (1973), scholars have documented persistent labor market disparities across gender and racial lines, motivating decades of policy interventions aimed at promoting equal employment opportunity.

These policy interventions utilize distinct approaches with varying enforcement methods, from direct numerical requirements to transparency-based oversight mechanisms. Norway's 2003 gender quota mandating 40 percent female board representation drove rapid increases in women's corporate board participation and improved market valuations, illustrating how enforceable targets can reshape organizational composition (Ahern and Dittmar, 2012). However, such direct mandates can also trigger unintended displacement effects, as demonstrated by California's SB 826, which produced immediate increases in female directors but also prompted backlash in non-regulated hiring channels (Bian et al. 2025). More subtly, Adams and Ferreira (2009) find that while female directors improve board monitoring practices, the average effect of gender diversity on firm performance is negative, particularly in well-governed firms. This suggests that mandatory diversity requirements can lead to "over-monitoring" that reduces shareholder value, demonstrating how well-intentioned mandates may inadvertently harm organizational performance.

Alternatively, disclosure-based approaches emphasize transparency and public accountability to drive behavioral change. Research on diversity reporting shows that organizations exhibit greater compositional improvements when they publish concrete targets rather than vague commitments (Cai et al. 2024). However, disclosure-only policies face risks of superficial compliance, with some organizations engaging in "diversity washing" to meet reporting obligations without substantively altering workforce demographics (Baker et al. 2024).

NYC's Bias Audit Law combines elements of both approaches, requiring independent third-party evaluation while mandating public disclosure of results. This hybrid regulatory design mirrors successful aspects of earlier interventions: like board gender quotas, it imposes binding compliance requirements. Like diversity disclosure mandates, it emphasizes transparency. The law's focus on algorithmic systems presents novel challenges but follows established patterns where regulatory oversight drives compositional changes when audit findings translate into corrective actions rather than cosmetic adjustments, providing the theoretical foundation for our empirical investigation of whether mandatory bias audits can effectively alter hiring practices while avoiding the displacement effects observed in earlier interventions.

3.3 Third-Party Auditing Theory and Effectiveness

The theoretical foundation for third-party auditing rests on information economics and agency theory, which explains how independent verification can mitigate information asymmetries and align incentives between principals and agents. In financial markets, decades of research demonstrate that external audits improve information quality, reduce earnings management, and enhance investor confidence (Francis 2004; DeFond and Zhang, 2014). The monitoring hypothesis suggests that audit scrutiny deters opportunistic behavior by increasing the likelihood of detection and reputational consequences (Watts and Zimmerman, 1983).

The expansion of auditing beyond financial statements has recently gained momentum, with assurance services emerging in several domains, including ESG metrics and decentralized protocols. Recent studies show that ESG assurance is becoming increasingly prevalent in the U.S. (Gipper et al., 2025). Similarly, the DeFi ecosystem has developed smart contract audits—voluntary code integrity assessments performed by specialized technical firms (Bourveau et al., 2024). Building on these developments, our study extends the role of audits into the algorithmic domain, where specialized AI consulting firms and technology companies currently lead the market for bias audits, reflecting the technical expertise required—similar to how non-traditional audit providers with specialized knowledge dominate ESG and DeFi assurance markets.

The application of auditing principles to algorithmic systems has generated substantial scholarly interest as policymakers seek mechanisms to address AI bias in employment decisions. The theoretical foundation for algorithmic auditing draws directly from established audit theory, where independent third-party verification serves to reduce information asymmetries and align organizational behavior with stakeholder interests (Francis 2004; DeFond and Zhang 2014). This framework suggests that systematic audits of hiring algorithms can help diagnose hidden group-based selection disparities. Kim (2017) argues that auditing increases firms' awareness of algorithmic impacts on hiring decisions, encouraging corrective actions and deterring discriminatory practices. This perspective aligns with broader audit theory, suggesting that third-party oversight improves organizational behavior even when audit methodologies are imperfect.

However, the effectiveness of bias audits in practice remains contested in the literature. Some research highlights fundamental limitations in current audit approaches. Kroll et al. (2017) and Leicht-Deobald et al. (2019) demonstrate that algorithmic audits often treat decision processes as black boxes, documenting disparate outcomes without identifying root causes or assessing substantive significance. This approach provides limited remediation guidance and may enable superficial adjustments that preserve underlying biases while creating an appearance of compliance. Additional concerns emerge regarding unintended consequences of audit mandates. Moschella (2022) and Celiktutan et al. (2024) warn that regulatory pressure may prompt employers to abandon AI screening in favor of human decision-makers, potentially reintroducing unconscious stereotypes that are more difficult to detect systematically. This substitution effect could undermine the policy objective of reducing employment discrimination.

The tension between these perspectives reflects broader challenges in AI governance, where technical complexity intersects with legal and social objectives. NYC's Bias Audit Law offers an opportunity to evaluate these competing predictions about algorithmic audit effectiveness in a real-world regulatory context. While prior research has largely relied on theoretical arguments and case studies, the mandatory nature of this regulation and its clear implementation timeline enable plausibly causal identification of bias audit effects on hiring outcomes.

3.4 Hypothesis Development

Drawing from the theoretical frameworks above, we develop hypotheses that address two key empirical questions: First, do bias audit requirements meaningfully alter hiring composition in ways consistent with reducing algorithmic discrimination? And under what conditions are these effects most pronounced? Second, what are the implementation costs of audit compliance?

Hypothesis 1 (Direct Effects on Hiring Composition): We test the null hypothesis that bias audit requirements have no effect on hiring composition. If bias audits merely treat algorithmic systems as "black boxes" without identifying root causes of discrimination, prompt organizations to abandon AI screening for potentially more biased human decision-making, or generate only cosmetic adjustments that preserve underlying biases, we should observe no systematic changes in hiring patterns following audit implementation. However, if companies undergo bias audits that identify algorithmic discrimination and remediate the underlying AI issues, or they incorporate more deliberate and gender-conscious human screening processes, we should observe systematic changes in hiring patterns following audit implementation. The monitoring theory suggests that third-party oversight creates incentives for organizations to modify practices that generate disparate outcomes, either through audit-driven remediation of algorithmic bias or through more deliberate human oversight of recruiting processes, even if underlying preferences remain unchanged. We therefore examine whether bias audits manifest as increased representation of previously underrepresented groups. We focus on understanding the share of female and minority hires, given anecdotal evidence that AI recruiting tools can mimic prior biases, even unintentionally, for these categories.

H1a: *Audit requirements have no effect on the share of male hires.*

H1b: *Audit requirements have no effect on the share of minority hires.*

We note that these hypotheses, in their alternative forms, jointly test the existence of disparate impact during the pre-period, particularly for companies using AI tools, and a reduction in disparate impact following the Bias Audit Law. In additional tests, we also explore two mechanisms by which audit requirements can lead to changes in hiring. First, we focus on firms that obtain bias audits. Second, we focus on whether firms revert to more manual oversight of AI recruiting processes. We also focus on firms with severe hiring imbalances, as they may employ AI recruiting tools that inadvertently exacerbated such imbalances during the pre-period.

Hypothesis 2 (Implementation Costs and Efficiency Trade-offs): Compliance with bias audit requirements imposes both direct costs (commissioning audits, coordinating with third parties) and indirect costs (expanding candidate searches, modifying screening processes). These compliance burdens should manifest in longer hiring timelines as organizations navigate new procedural requirements. In addition, firms might increase the number of HR staff in reaction to the law and even shift recruiting away from NYC. We propose the following hypothesis.

H2: *Audit requirements have no effect on hiring timelines, HR staff, and total hiring volume.*

4. Research Design and Data

4.1 Research design

To examine the effect of the Bias Audit Law on hiring outcomes (H1a and H1b), we use a difference-in-differences (DiD) approach that exploits geographic variation in the law's implementation. We compare treatment offices located in NYC and control offices of the same company located outside NYC.[11] The control sample is restricted to firms that have at least one office in NYC to ensure within-firm comparability. This restriction ensures that both treatment

[11] Because Revelio Labs provides location data only at the state and metro area levels, we define the New York area as both the state being New York and the metro area being New York. This is reasonable because NYC comprises 70% of the metro area, functions as a unified labor market with significant commuting, and hosts most large employers subject to audit requirements. Additionally, if the same company appears in different states or metro areas, we treat each as a separate office.

and control offices belong to firms that are subject to the same corporate policies, culture, and strategic decisions, with the key difference being their exposure to the Bias Audit Law.

We focus on a nine-quarter window centered on the quarter in which the regulation was enforced (2023Q3). Figure 2 illustrates the research design. Panel A shows the event window timeline, and Panel B depicts the identification of treatment and control offices. Our office-quarter level regression specification is stated in equation (1):

$$Y_{i,t} = \beta_0 + \beta_1 Treat_i * Post_t + \gamma X_{i,t} + u_i + v_t + \varepsilon_{i,t} \quad (1)$$

where *i* indexes offices and *t* quarters. $Y_{i,t}$ is hiring diversity, comprised of two measures: *MaleProp* and *WhiteProp*, the proportions of male and white employees, respectively, hired in office *i* during quarter *t*. These variables allow us to examine how gender and racial compositions of new hires respond to the regulation.

[Insert Figure 2 about here]

The key explanatory variable is the interaction between *Treat* and *Post*. The indicator variable, *Treat*, equals one for offices located in NYC. The other indicator variable, *Post*, equals one for quarters on or after 2023 Q3, when the Bias Audit Law became effective. The coefficient of interest, β_1, captures the differential change in hiring outcomes for NYC offices relative to control offices following the implementation of the Bias Audit Law. We control for *LogEmpNum*, the natural logarithm of the number of employees in office *i* during quarter *t*. To mitigate the influence of outliers, we winsorize all continuous variables at the 1% and 99% levels. Appendix Table A provides detailed variable definitions.

We include office and quarter fixed effects to control for time-invariant characteristics such as long-standing HR practices and to absorb aggregate shocks and seasonal patterns such as nationwide recruitment cycles and hiring surges. Thus, the stand-alone coefficients on *Treat* and

Post are subsumed by this fixed effect structure. Standard errors are clustered at the office level to account for potential serial correlation in error terms within offices over time.

4.2 Data and sample selection

We construct the sample using data from several sources. The primary data source is from Revelio Labs, which collects and aggregates all publicly available professional profiles and job postings.[12] The Revelio Individual database provides an extensive range of variables for each individual, including demographic details such as gender and ethnicity, educational backgrounds such as academic degrees and start and end dates, and employment histories such as location, seniority level, and company industry. The Revelio Workforce Dynamics database includes the total number of employees at each office on a monthly basis. The Revelio Job Postings database contains job postings sourced from aggregator sites and company websites. We additionally use a technology database from TheirStack that documents technologies adopted by each firm, which we use to identify whether a firm has adopted audited AI hiring algorithms.[13]

To identify which firms had their algorithm audited, we consider two groups. The first group includes firms that have directly issued a bias audit report, which we identify by manually scraping each company's official website to check for publicly available bias audit reports, and supplement this with the ACLU-maintained crowd-sourced tracker.[14] The second group are firms that indirectly had their algorithm audited by adopting audited AI hiring algorithms covered in TheirStack's technology database, whose vendors have publicly posted bias audit reports. More

[12] We acknowledge that Revelio Labs primarily captures professional and white-collar workers, which limits the generalizability of our findings to this segment of the labor market rather than to blue-collar or service workers. However, this data characteristic also helps address concerns about confounding events such as strikes and union organizing activities, as professional and white-collar workers are substantially less unionized.
[13] TheirStack's technographic data maps 32 thousand technologies across 12 million companies worldwide collected from sources such as job postings and company websites.
[14] Available at https://github.com/aclu-national/tracking-ll144-bias-audits

details on the identification of both groups are provided in Section 6.1.

Our main sample draws from these databases, focusing on US offices from 2022 Q2 to 2024 Q3. For both the Revelio Individual and Job Postings databases, we aggregate the data at the office-quarter level. The Workforce Dynamics data is aggregated quarterly and then merged with our sample for the respective hiring diversity and recruitment efficiency regressions. We then merge with the technology database to identify the technology tools adopted by each firm.

We require each office to have both pre- and post-treatment periods. Additionally, each treated office must have a corresponding control office located outside the NYC area but within the same company. Our final sample comprises 2,522 companies, 91,600 offices, and 477,575 office-quarter observations for the main hiring regression.[15] Among these, 1.1% of companies directly issued a bias audit report and 38.1% had their algorithm indirectly audited by adopting one of the 11 audited vendor algorithms. Sample sizes vary slightly across different tests due to the use of different dependent variables.

For office-level performance analyses, we examine changes in financial analyst forecast accuracy and employee satisfaction. Financial analyst forecast data are obtained from I/B/E/S, with analyst office locations identified by scraping analyst reports from Mergent Online. Employee satisfaction data are obtained from the Revelio sentiment database.

5. Main Results

5.1 Descriptive statistics

Table 1 Panel A reports summary statistics for outcome and control variables. For the full

[15] Our initial sample from the aggregated Revelio Individual database contains 150,727 companies and 645,460 offices. After merging this data with the Workforce Dynamics database to obtain the control variable for the total number of employees, the sample size is significantly reduced. However, our results remain robust when using the large initial sample.

sample, the average proportion of male hires (*MaleProp*) is 56.3%, and the average proportion of white hires (*WhiteProp*) is 69.9%. The table also includes the proportions of different age groups, as well as the proportion of White male hires and the proportion of Non-White male hires. The average log job vacancy duration (*LogDuration*) is 3.539, corresponding to about 34 days.[16] The average ratio of net HR staff hires (*HRHireRate*), and total staff hires (*HireRate*) in an office are -0.024 and 0.081 respectively. The average log number of employees (*LogEmpNum*) is 4.811, indicating a typical office size of about 123 employees.

These variables exhibit substantial variation across offices and quarters. *MaleProp* has a standard deviation of 0.396, with a 0.083-1.000 interquartile range, indicating that some offices hire predominantly female workers while others hire exclusively male workers. *WhiteProp* also varies considerably, though more moderately than *MaleProp* (standard deviation = 0.362, interquartile range: 0.500-1.000), suggesting greater gender imbalance relative to ethnicity / racial imbalance across offices. There is substantial heterogeneity in hiring activity across offices, with interquartile *HRHireRate* ranging from −0.061 to 0 and *HireRate* ranging from −0.315 to 0.078. The log variables also display considerable variation, with *LogDuration* ranging from 3.258 to 3.885 and *LogEmpNum* from 3.524 to 6.057across the interquartile range.

Panel B reports treatment-control comparisons in both the pre- and post-periods for our key dependent variables, where treatment offices are NYC offices and control offices are non-NYC offices within the same firm. The univariate difference-in-differences show a significant decline in the proportion of male hires and a significant increase in job vacancy duration for NYC offices relative to control offices. Panel C and Panel D repeat this comparison, partitioning treated NYC offices into those which had their algorithm audited either directly or indirectly (i.e., *Report*

[16] The average for the unlogged *Duration* is higher, at 40 days, consistent with some skew existing in the variable.

firms), and those belonging to *NonReport* firms, comparing each group against the same control group. Both groups show a significant decline (increase) in male hiring share (job vacancy duration). However, the decline in male hiring share is larger for firms with audited algorithms, consistent with audits generating stronger effects. The significant decline and longer vacancy duration among firms without audits suggests that some firms revert to more time-consuming but potentially more gender-conscious human screening.

[Insert Table 1 about here]

Panels B to D also provide additional insights. While the proportion of male hiring is larger in NYC than outside of NYC during the pre-period, it is almost the same during the post period. In contrast, job vacancy durations, almost the same during the pre-period, substantially increase for NYC offices during the post period. Finally, the NYC male hiring share, larger during the pre-period by 1.5 percentage points for audited vs. non audited offices, becomes almost the same in the post-period. We confirm this univariate result in multivariate analyses that control for office size and industries and are reported in Table A1 of the online appendix. This result provides some evidence that companies relying on AI tools may have exhibited greater bias against female applicants in the pre-period relative to companies using more manual processes, although we caution against deriving strong conclusions from this analysis.[17]

5.2 Bias Audit and Hiring Diversity

Panel A of Table 2 reports the results of equation (1), with columns (1) and (2) [(3) and (4)] for the proportion of male (white) hires. For each outcome, we report specifications without (columns 1 and 3) and with the office size control variable (columns 2 and 4).

---

[17] Specifically, our study does not need to rely on the assumption that AI recruiting tools are more biased than manual human processes. Just on the fact that AI tools can exhibit some bias that ultimately gets corrected through the NYC Bias Audit Law. Our empirical evidence is collectively consistent with bias existing during the pre-period.

The coefficients on the interaction term *Treat*×*Post* in the first two columns of Table 2 Panel A are statistically negative and economically meaningful (−0.018, t = −3.397 in column 1; −0.018, t = −3.418 in column 2). This suggests that NYC offices relatively reduce their proportion of male hires on average by almost 2 percentage points following the implementation of the Bias Audit Law, representing a 3.2% reduction from the mean share of male hires. In contrast, the coefficients for *WhiteProp* are insignificant across both specifications. These results suggest an average decline in male hiring after the implementation of the Bias Audit Law, while no evidence is found for white hiring, thus providing evidence to reject H1a but not H1b.

One possible explanation for this difference is the disparity in algorithmic identification accuracy. Both gender and race information in the Revelio Labs dataset and in resume screenings are inferred from individuals' names. However, name-based algorithms generally predict gender with substantially higher accuracy than race (Lockhart et al. 2023). Moreover, evidence from the Amazon case suggests that bias against women in AI-based hiring algorithms is common.

The coefficient on the control variable *LogEmpNum* is consistently negative and statistically significant, indicating that larger offices tend to have more diverse hiring patterns. This finding aligns with previous literature documenting that larger organizations typically maintain more formalized diversity practices (Oliveira and Zhang 2022).

[Insert Table 2 about here]

A key underlying assumption is that, in the absence of the Bias Audit Law, the trends for the treatment and control groups would have remained parallel. While this assumption cannot be directly tested, we can still assess whether the pre-treatment trends are parallel. We re-estimate Equation (1) by replacing the *Post* variable with separate indicators for each quarter before and after the regulation. The benchmark quarter is 2023Q2, which immediately precedes enforcement of Bias Audit Regulations. The dynamic results for the male hiring proportions are shown in Figure 3, which displays point estimates and 90% confidence intervals. Point estimates for all quarters before 2023 Q3 are statistically indistinguishable from zero, suggesting no systematic differences between treatment and control offices prior to the law's enforcement. In contrast, we observe a relatively significant decrease in male hiring post-enforcement.[18]

[Insert Figure 3 about here]

5.3 Bias Audits and Hiring Diversity: White Males and Non-White Males

We next examine how the Bias Audit Law affects male hiring across racial lines. This allows us to determine whether the reduction in male hiring is concentrated among specific groups or represents a broad shift toward greater gender diversity across all racial categories.

The corresponding results are reported in Panel B of Table 2. Column (1) shows the proportion of white-male hires, and column (2) shows the proportion of non-white male hires. The coefficients on *Treat*×*Post* are significantly negative in both columns (−0.010, t = −1.846 for white males; −0.008, t = −1.922 for non-white males), indicating that the reduction in male hiring occurs across both white-male and non-white-male cohorts. This finding is consistent with the law's broader aim of addressing multiple dimensions of bias, rather than targeting a single demographic group.

5.4 Addressing Confounding Events with a Placebo Test: Bias Audits and Age Diversity in Hiring

A concern is that our results may reflect broader contemporaneous changes in NYC's labor market, such as the city's strong diversity culture and initiatives or other labor market shifts, rather than the specific impact of bias audits. To strengthen the causal interpretation of our

[18] A joint F-test of the pre-period coefficients fails to reject the null hypothesis that all pre-period coefficients are jointly equal to zero, further strengthening support for the parallel trends assumption.

findings and rule out alternative explanations, we conduct placebo tests examining age diversity in hiring patterns. The Bias Audit Law requires audits for race and gender categories but does not specifically require addressing age-based discrimination (e.g., Hickock 2022). This institutional feature provides a natural falsification test: if our observed effects on gender hiring truly result from the Bias Audit Law rather than confounding factors, we should observe no systematic changes in age-related hiring patterns following the law.

We construct age-based hiring proportions using four cohorts based on demographic classifications employed by the U.S. Bureau of Labor Statistics: employees under 25 (*PropAgeUnder25*), ages 25-34 (*PropAge25To34*), ages 35-44 (*PropAge35To44*), ages 45 and above (*PropAge45Plus*). The results, presented in Table 3, show no significant changes in hiring shares across any age group following implementation of the Bias Audit Law.[19] This finding supports the conclusion that the changes in gender diversity are not confounded by other contemporaneous policies, such as diversity campaigns, pay transparency mandates, shifts in labor supply following the pandemic, or the candidate notification requirement under the law.[20]

[Insert Table 3 about here]

6. Mechanism Tests and Additional Analyses

Having established that the Bias Audit Law reduces male hiring, we explore plausible underlying mechanisms and boundary conditions of these effects. Two plausible mechanisms may explain our results. First, bias audits can directly identify and remediate algorithmic

[19] The results remain similar (i.e., insignificant) if we partition NYC offices between reporters and non-reporters, as defined later in the paper.

[20] If candidate notifications about the use of AEDT in hiring were the major driver of our results, we would expect shifts in age composition, since age discrimination is documented in AEDT use, such as iTutorGroup's lawsuit for using AI to exclude older applicants (Wiessner 2023). The absence of such changes supports attributing the observed effects to the bias audit itself.

practices that produce disparate impact. This requires firms to undergo a bias audit. Second, firms may respond to the law by scaling back AI recruiting and increasing human oversight with an explicit focus on reducing bias, potentially circumventing the need for a formal audit. We conduct a series of analyses to examine both mechanisms. We also assess potential costs and benefits of the Bias Audit Law by analyzing its impact on HR hiring, recruitment efficiency, and firm performance following its enforcement, and conduct several additional tests.

6.1 Cross-Sectional Test: Audit-Reporters and Non-Reporters

The first mechanism is that the Bias Audit Law leads firms to identify and remediate algorithmic practices that produce disparate impact through explicit audit compliance. To examine this mechanism, we identify firms in our sample that had their algorithm audited, either directly or indirectly. We first identify firms that have directly issued a bias audit report. Since the law requires public disclosure of audit results on company websites for at least six months, we scrape each company's official website to find the corresponding bias audit report. As firms may remove reports after the disclosure period, we supplement our scraping with the ACLU-maintained crowd-sourced tracker. We find that direct compliance, where a third party conducts a bias audit, is limited to approximately 1% of firms, consistent with Wright et al. (2024).

However, firms that indirectly had their algorithm audited through their vendor are substantially more prevalent. Many firms adopt off-the-shelf AI recruiting solutions from data vendors, and AI hiring algorithm vendors proactively commission independent “aggregate” bias audits and publicly post the results to market their tools as unbiased. When those vendors conduct their own bias audits, their customers are indirectly covered by the same audit results. Through the bias audit reports we scraped, we identify 11 AI hiring algorithms that have been independently audited and are also covered in TheirStack’s technology database. Next, we identify which firms in our sample adopted at least one of these algorithms. Overall, we estimate that 39% of firms in our sample had their algorithm audited, either directly or indirectly.

If this mechanism drives the results, we would expect the reduction in male hiring to be more pronounced among firms that comply with the law. To test this prediction, we construct two indicators (1) *Report*, which equals one if a firm has issued a bias audit report, either directly or indirectly; and (2) *NonReport*, which equals one for firms that are not classified as *Report* firms. *NonReport* pools firms that rely primarily on human screening and are not subject to the Bias Audit Law, firms whose audit reports we may have missed, and non-compliant firms. We then estimate triple-difference specifications by interacting these indicators with *Treat×Post*.

Table 4 Panel A presents the results. The coefficient on *Treat×Post×Report* is significantly negative with a larger magnitude (-0.029) than the negative but insignificant coefficient on *Treat×Post×NonReport* (-0.01), confirmed using an F-test of differences. The concentrated and stronger effects among firms that had their algorithm audited suggest that the auditing process itself plays a primary role, rather than public disclosure or general regulatory pressure, as most *Report* firms are indirect reporters who made no public disclosure.

To further support the mechanism that the auditing process can help identify and remediate biased algorithmic practices, we test whether the effect is stronger for algorithms that were more gender-imbalanced before the law, as more biased algorithms offer greater room for the audit to identify and correct disparate impact. We construct *More(Less)Biased,* which equals one for algorithms whose adopting firms exhibited higher (lower) mean male hiring rates before the law, respectively, and interact these with *Treat×Post×Report*. As shown in Table 4 Panel B, both *Treat×Post×Report×MoreBiasedV* and *Treat×Post×Report×LessBiasedV* are significantly negative, but the former is significantly more negative than the latter. This indicates that the effect is more pronounced among offices that adopt algorithms trained on more gender-imbalanced populations. Taken together, these results are consistent with bias audits leading to the identification and remediation of algorithmic practices that produce disparate impact.

We then explore potential spillover effects to non-NYC offices for firms that directly or indirectly obtain a bias audit. Panel C replaces *Treat×Post* in equation (1) with *NY office×Post×Report*, *Non-NY office×Post×Report*, and *NY Office×Post×NonReport*. where the benchmark group is comprised of non-NYC offices of non-audited firms. As shown in Panel C of Table 4, *Non-NY office×Post×Report* is also significantly negative, though its magnitude is significantly smaller than that of *NY office×Post×Report*. These results suggest that audit-driven remediation in NYC partially extends to non-NYC offices. However, the smaller reduction in male hiring outside NYC is consistent with firms not uniformly adopting the same hiring processes across all offices or taking additional corrective action in their NYC offices. We focus on some of these actions below.

[Insert Table 4 about here]

6.2 Cross-Sectional Test: Heterogeneous Effects by Pre-Law HR Staff Share

The second mechanism is that firms respond to the law by scaling back AI recruiting and increasing human oversight with an explicit focus on reducing potential biases. We explore this mechanism by studying the role of HR capacity in monitoring AI algorithms and examining whether firms actively expand their capacity in response to the law.

We first examine whether the reduction in male hiring is stronger in offices with more limited HR support, as these offices likely rely more heavily on AI systems and are therefore more susceptible to algorithmic bias when human oversight is limited (Skitka et al. 1999; Wilson et al. 2025). To test this prediction, we construct two indicators (1) *LowHR*, which equals one for offices whose pre-law HR staff share is below the sample median, and (2) *HighHR*, which equals one for offices whose pre-law HR staff share is above the sample median. We then estimate the specifications by interacting these indicators with *Treat*×*Post*. As shown in Table 5 Panel A, *Treat*×*Post*×*LowHR* is significantly negative with larger magnitude while *Treat*×*Post*×*HighHR* is insignificant, and the two coefficients are significantly different from each other. This suggests that the reduction in male hiring is concentrated in offices with more limited HR support.

We further examine the joint heterogeneous effect of HR support and audit compliance by decomposing *Treat*×*Post* in equation (1) into four interaction terms based on HR support and audit status. As shown in Table 5 Panel B, only *Treat*×*Post*×*Report*×*LowHR* is significantly negative and significantly different from the other three coefficients, indicating that the reduction in male hiring is strongest and concentrated among offices that had their algorithm audited and had limited HR support, where reliance on AI is high and human oversight is limited.

To further examine whether offices respond to the Bias Audit Law by enhancing human oversight, we construct *HRHireRate*, the net HR staff hires, measured as the change in HR headcount from year t−1 to year t scaled by HR headcount in year t−1, as the dependent variable. Column (1) of Table 5 Panel C reports the results of the main difference-in-differences specification, while column (2) further decomposes the main interaction term to distinguish between audit-reporters and non-reporter firms. *Treat*×*Post* is significantly positive in column (1), and both *Treat*×*Post*×*Report* and *Treat*×*Post*×*NonReport* are significantly positive in column (2), suggesting that offices increase hiring HR staff following the law, regardless of whether their algorithm was audited. Offices obtaining bias audits may expand human oversight to complement audit-driven remediation, while offices without audit evidence may revert to more human-intensive screening processes to circumvent the bias audit requirement.

[Insert Table 5 about here]

We also examine the dynamic effect on HR staff net hires in Figure 4. Point estimates for all quarters before 2023 Q3 are statistically indistinguishable from zero, suggesting no systematic differences between treatment and control offices prior to the law's enforcement. A significant increase in HR staff net hires emerges after the enforcement of the Bias Audit Law. This result is consistent with firms resorting to more manual processes in response to the Bias Audit Law.

[Insert Figure 4 about here]

Taken together, these results are consistent with the second mechanism, whereby firms increase human oversight in response to the law, either to complement audit-driven remediation of their AI systems or to circumvent the bias audit requirement entirely by reverting to human hiring processes.

6.3 Bias Audit and Job Vacancy Duration

This section examines an additional cost of Bias Audit Laws beyond the hiring of HR staff (H2), the time offices spend filling open positions. Affected offices may adjust their recruitment practices in response to the Bias Audit Law. For example, firms could devote more time to administrative compliance for third-party audits, expand candidate pools to increase diversity, add supplementary screening steps, or conduct more thorough evaluations before the next audit cycle. Thus, increased hiring diversity may come at the cost of reduced recruitment efficiency.

The results are presented in Table 6 Panel A. Column (1) reports OLS estimates with *LogDuration*, the natural logarithm of the average number of days to fill a vacancy in a quarter, as the dependent variable. Column (2) shows Poisson estimates with *Duration*, the average number of days to fill a vacancy in a quarter, as the dependent variable. The coefficients on *Treat*×*Post* are significantly positive in both specifications (0.189, t=18.535 for OLS regression;

0.200, t=21.134 for Poisson regression), indicating that NYC offices experience a 20%-22% increase in average job vacancy duration relative to control offices. Given the baseline average vacancy duration of 34 days, the increase translates to approximately 7 additional days per hire.

In addition, Table 6 Panel B shows that both *Treat×Post×Report* and *Treat×Post×NonReport* are significantly positive with similar magnitude, suggesting that the increase in job vacancy duration is present regardless of audit status. For offices with audited algorithms, the longer vacancy duration may reflect expanded candidate searches and modifications to screening processes following the audit. For non-reporter offices, the increase is consistent with a reversion to more time-consuming human screening.

[Insert Table 6 about here]

We also examine the dynamic effect of the Bias Audit Law on job vacancy duration. As shown in Figure 5, no significant differences in job vacancy duration are visible between the treatment and control groups in the pre-adoption period, supporting the parallel trend assumption. However, post-adoption, we observe a significant increase.

[Insert Figure 5 about here]

Overall, these findings suggest that while the Bias Audit Law increases female hiring, it also increases HR staffing and decreases recruitment efficiency by extending job vacancy duration. We caveat, though, that because this analysis uses Revelio Labs data, it primarily applies to professional and white-collar professions.

6.4 Bias Audit and Total Hires

Another potential cost is that companies may reduce recruitment for their NYC offices and shift hiring to other locations. To investigate this possibility, we focus on the office-level net hires (*HireRate*), measured by the change in the number of employees divided by prior-period number of employees. Table 7 presents the results. The coefficients on *Treat×Post* are insignificant in Panel A, suggesting no overall reduction in hiring volume at NYC offices following the law. We further examine whether hiring rates differ between reporter and non-reporter offices. Both *Treat* × *Post* × *Report* and *Treat* × *Post* × *NonReport* are insignificant in Panel B, indicating that the absence of a hiring reduction holds regardless of audit status.

We next explore whether hiring is displaced to neighboring offices. Panel C augments the baseline specification with an additional interaction term, *NeighborOffice* × *Post*, where *NeighborOffice* equals one for offices located within the New York MSA but outside New York State. This interaction is also insignificant, providing no evidence of hiring displacement to nearby offices. Taken together, these results indicate no detectable decrease in hiring volume or geographic shifting of hiring activity in response to the law.

[Insert Table 7 about here]

6.5 Heterogeneous Effects by Pre-Law Male Dominance

Having established that the law reduces male hiring and increases job vacancy duration, and the co-existence of two mechanisms, audit-driven remediation and a shift towards more human screening, we examine whether the magnitude of these effects varies with the degree of pre-existing gender imbalance in firms' hiring environments. Firms in more male-dominated environments may have paid less attention to gender diversity before the law, and the regulatory pressure may prompt more substantial adjustments to their hiring practices. Further, AI-based screening models may put too much emphasis on the characteristics of overrepresented groups, leading to inadvertent screening out of qualified candidates from other groups, similar to the Amazon 2014 AI recruiting experiment (Dastin 2018; Goodman 2018). We thus expect the reduction in male hiring to be more pronounced in firms with more severe pre-audit gender imbalances. We construct (1) *High(Low)MaleFirm*, a measure of pre-law male hiring imbalance which equals one if a firm's pre-law male hire share is above (below) the sample median, and (2) *High(Low)MaleInd*, a measure of industry-level gender composition which equals one if a firm is in a male(female)-dominated industry, defined as industries where the female employment share is below(above) the overall average across all industries. We then estimate the specifications by interacting these indicators with *Treat*×*Post*.

Table 8 presents results consistent with our expectations. Across all three outcomes, the effects are generally stronger for firms and industries with more severe pre-existing gender imbalances. The reduction in male hiring and the expansion of human oversight are concentrated among high-male firms and industries, while low-male firms show smaller and insignificant effects, indicating that firms in male-dominated environments are more likely to respond to the law by actively increasing their HR capacity and increasing their hiring diversity. The increase in job vacancy duration is present across all groups but is larger for high-male firms and industries, suggesting that while the law imposes broad compliance costs, the adjustments are more substantial where gender imbalance was more severe.

Table 8 Panel B further examines these effects jointly by gender imbalance and audit status. For male-dominated firms and industries, both audit-reporter and non-reporter offices show a reduction in male hiring, an expansion of human oversight, and longer job vacancy durations, with no statistically significant difference between the two groups. These results further support the co-existence of two mechanisms under the Bias Audit Law. Even for firms that do not commission a bias audit, the law can lead them to revert to human screening processes. Male-dominated environments represent one such instance where this shift is consequential enough to produce hiring outcomes comparable to those driven by the auditing process itself.

[Insert Table 8 about here]

6.6 Bias Audit and Subsequent Firm Performance: Evidence from Brokerage Firms

In this section, we further investigate whether the hiring changes brought by the Bias Audit Law translate into office-level performance improvement. We focus on brokerage firms because financial analysts are typically based in a specific office, allowing analyst forecast accuracy to serve as a direct measure of performance (e.g., Clement 1999; Mikhail et al. 1999; Loh and Mian 2006). First, we confirm that the hiring effects in our full sample extend to brokerage firms. Next, we analyze how financial analysts' output quality, measured by analyst EPS forecast error, changes following the implementation of the Bias Audit Law.

The results are presented in Table 9. Panel A confirms that brokerage NYC offices experience a reduction in male hiring and increases in HR staff hiring and job vacancy duration, consistent with the full sample results. Panel B reports the summary statistics for variables in the analyst performance analysis. Panel C reports the effect on financial analyst performance, using analyst forecast error as the dependent variable. Column (1) [(2)] focuses on the overall effect (on heterogeneity between reporter and non-reporter offices). The coefficient on *Treat*×*Post* is statistically negative in column (1), consistent with a positive effect of the Bias Audit Law on performance of analysts at NYC office. Further, in column (2), only *Treat* × *Post* × *Report* is significantly negative while *Treat* × *Post* × *NonReport* is insignificant, and the F-test confirms the two coefficients are significantly different from each other, suggesting the improvement in financial forecasting ability is concentrated and stronger among offices that obtain bias audits.

[Insert Table 9 about here]

Overall, our results provide some evidence that discipline introduced by Bias Audits leads to improvements in brokerage firms' performance. Thus, these results help answer whether companies truly address fairness issues in recruiting or make arbitrary changes to their processes that do not address fairness issues in recruiting. Under the former scenario, we would expect the newly hired employees to be of higher quality on average, thereby resulting in performance improvements. Under the latter, we would expect unchanged or even decreased performance.

6.7 Bias Audits and Employee Satisfaction: Evidence from Employee Reviews

To further assess whether bias audits requirements address actual workplace bias and explore broader social effects, we examine how new hires' perceptions of their workplace change following the law. We use employee satisfaction ratings from Revelio Labs' sentiment database, restricting the sample to employees with less than one year of tenure to focus on new hires. Table 10 reports the results. Panel A presents the baseline treatment effects across six rating categories: *Overall Rating*, *Work-Life Balance*, *Senior Management*, *Diversity and Inclusion*, *Career Opportunities*, and *Culture and Values*. Panel B compares the effects between audit reporters and non-reporters.

In Panel A, the coefficients *Treat* × *Post* are significantly positive for *Overall Rating*, *Diversity and Inclusion*, and *Career Opportunities*, consistent with the bias audit law's focus on fairness in recruitment and advancement. The remaining dimensions, such as work-life balance, and management styles, are insignificant, which is expected as these dimensions are not directly targeted by the bias audit law and can therefore serve as placebo checks.

In Panel B, the coefficients *Treat* × *Post* × *Report* are significantly positive for *Overall Rating*, *Diversity and Inclusion*, and *Culture and Values*, whereas *Treat* × *Post* × *NonReport* is insignificant across all six dimensions. These results suggest that the improvement in diversity-related ratings is concentrated among offices that conducted bias audits. Collectively, the findings provide evidence that, beyond reducing disparate impact in hiring, the law may also help mitigate underlying biases in the workplace.

[Insert Table 10 about here]

6.8 Additional Analyses and Robustness Tests

We conduct several analyses that are reported in the online appendix. First, we focus on job seniority and find in Online Appendix Table A2 that our results hold for junior and mid-level positions, where AI recruiting systems are more likely used, but not for senior positions, where AI recruiting systems are less likely to be used. This latter result provides another meaningful placebo test. The result on mid-level positions also helps alleviate concerns that NYC firms respond more to recent national trends such as unionization drives, as these trends tend to disproportionately apply to lower- or junior-level positions.

Second, we conduct subsample analyses based on firm public status, and find in Online Appendix Table A3 that our results apply for both public and private firms. Given that public firms are subject to greater reputation and regulatory risks relative to private firms, these results help rule out that such forces primarily drive our results.

Third, we conduct an additional analysis focusing on non-white subgroups in Table A4, and find some heterogenous effects on non-white subgroup hiring. This suggests that Bias Audit Laws may lead to a redistribution of hiring opportunities among non-white racial groups.

Fourth, we conduct several untabulated robustness tests. 1) To ensure our findings are not sensitive to quarterly aggregation, we re-estimate our models using office-month data. This higher-frequency specification provides additional temporal variation and addresses concerns that quarterly aggregation might mask important month-to-month dynamics or create artificial patterns in the data. 2) To account for the possibility that New York's strong commitment to diversity may distinguish it from other labor markets, we narrow our control group to offices of the same firms located in other cities within New York State. 3) We replace Year-Quarter fixed effects with Firm×Year-Quarter fixed effects. This specification provides a more stringent test by comparing treated and control offices within the same firm and time period, effectively controlling for any concurrent firm-level policies or shocks that might affect all offices simultaneously. 4) We address potential concerns about violations of the Stable Unit Treatment Value Assumption (SUTVA), which could occur if control offices are indirectly affected by the treatment. To mitigate this concern, we re-estimate the regression using an alternative control group that includes only offices from firms that have never had an office in NYC. Our main results hold across all of these robustness checks.[21]

7. Conclusion

This paper provides the first large-sample empirical evidence on the effects of Bias Audit Law on hiring outcomes. Using a Difference-in-Differences design, we find the law reduces male hiring share by almost 2 percentage points. We document two mechanisms. First, the effects are concentrated among firms that had their algorithm audited, either directly or indirectly through their vendor, suggesting that the audit process itself rather than public disclosure or general regulatory pressure drives the results. Second, firms broadly expand human oversight in response

[21] In untabulated analyses, we conduct several additional robustness tests. First, to ensure greater comparability between treated and control offices, we restrict the control sample to offices in large metropolitan areas with similar economic characteristics, such as Los Angeles, Chicago, Dallas, Houston, and Washington, D.C. Our main results remain robust. Second, we address concerns arising from the extended timeline between the law's introduction (February 2020) and enforcement (July 2023). During this three-year gap, firms may have gradually adjusted their practices or anticipated the requirement. We use a "clean" comparison period by restricting our sample to pre-introduction quarters (Q1-Q4 2019) and post-enforcement quarters, excluding the 2020-2022 interim period. Our results remain consistent, suggesting that transitory adjustments during implementation do not drive our findings. Then, we test for anticipation effects around the official announcement in Q4 2021. We examine two alternative windows: (1) Q1 2019-Q3 2021 versus Q4 2021-Q2 2023, and (2) Q4 2020-Q3 2021 versus Q4 2021-Q4 2022. We find no significant changes in male hiring around the announcement date, indicating no anticipation effects and confirming that our results reflect responses to actual enforcement.

to the law, as evidenced by significant increases in HR staff hiring across both audit-reporter and non-reporter offices. Our main effects are also stronger for firms with more gender-imbalanced algorithms and lower HR staff share. On the cost side, job vacancy duration increases by 20 to 22 percent, though total hiring volume and geographic distribution remain unchanged.

Using brokerage firms as a setting, we also find evidence of performance improvements in the form of greater analyst forecast accuracy at NYC brokerage offices, particularly among those whose algorithms were audited. We also find new hire satisfaction ratings improve in dimensions directly tied to fairness in recruitment.

Overall, our study offers timely insights for policymakers evaluating similar regulations globally and contributes to the growing literature on AI governance and the evolving role of audits in non-financial domains.

References

Adams, R. B., and D. Ferreira. 2009. Women in the boardroom and their impact on governance and performance. *Journal of Financial Economics* 94 (2): 291–309.

Adhikari, B. K., A. Agrawal, and J. Malm. 2019. Do women managers keep firms out of trouble? Evidence from corporate litigation and policies. *Journal of Accounting and Economics* 67 (1): 202–225.

Ahern, K. R., and A. K. Dittmar. 2012. The Changing of the Boards: The Impact on Firm Valuation of Mandated Female Board Representation *. *The Quarterly Journal of Economics* 127 (1): 137–197

Ahn, J., R. Hoitash, U. Hoitash, and E. Krause. 2025. Diversity and Career Trajectories: Evidence from LinkedIn Data on Race, Ethnicity, and Gender in Auditing. *The Accounting Review* 100 (4): 1–31.

An, J., D. Huang, C. Lin, and M. Tai. 2025. Measuring gender and racial biases in large language models: Intersectional evidence from automated resume evaluation. PNAS Nexus, 4(3)

Aobdia, D. 2019. Do practitioner assessments agree with academic proxies for audit quality? Evidence from PCAOB and internal inspections. *Journal of Accounting and Economics* 67(1): 144-174.

Aobdia, D., P. Choudhary, and N. Newberger. 2024. The economics of audit production: what matters for audit quality? An empirical analysis of the role of midlevel managers within the audit firm. *The Accounting Review* 99(2): 1-29.

Aobdia, D., G. Köchling, P. Limbach, and A. Yoon. 2025. Emissions Restatements After the SEC Climate Proposal: Evidence from Carbon Disclosure Project Filings. Working paper.

Arrow, K., Ashenfelter, O. and Rees, A., 1973. Discrimination in labor markets. *The Theory of Discrimination*, pp.3-33.

Association of International Certified Professional Accountants. AICPA. 2023. SOC for service organizations engagements.

Babina, T., A. Fedyk, A. He, and J. Hodson. 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151: 103745.

BABL AI. 2022. Public Comment on Proposed Rules Related to Local Law 144 (Automated Employment Decision Tools). New York City Department of Consumer and Worker Protection, June 3. Available at: https://www.nyc.gov/assets/dca/downloads/pdf/about/PublicComments-Proposed-Rules-Related-to-LocalLaw202of2019-LocalLaw1144of2021-LocalLaw37of2022.pdf.

Baker, A. C., D. F. Larcker, C. G. McClure, D. Saraph, and E. M. Watts. 2024. Diversity Washing. *Journal of Accounting Research* 62 (5): 1661–1709.

Barocas, S., and A. D. Selbst. 2016. Big Data's Disparate Impact. *California Law Review* 104: 671.

Becker, G. S. 1957. *The Economics of Discrimination*. University of Chicago Press.

Bian, B., J. Li, and K. Li. 2025. Does Mandating Women on Corporate Boards Backfire? SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Blankespoor, E., E. deHaan, and Q. Li. 2025. Generative AI in Financial Reporting. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Billings, M. B., A. Klein, and Y. C. Shi. 2022. Investors' response to the #MeToo movement:

does corporate culture matter? *Review of Accounting Studies* 27 (3): 897–937.

Bonfiglioli, A., R. Crinò, M. Filomena, and G. Gancia. 2025. Comparative Advantage in AI-Intensive Industries: Evidence from US Imports. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Bourveau, T., J. Brendel, and J. Schoenfeld. 2024. Decentralized Finance (DeFi) assurance: early evidence. *Review of Accounting Studies* 29 (3): 2209–2253.

Bourveau, T., X. Gao, and O.-K. Hope. 2025. The Impact of Disclosure on Diversity: Evidence from the Canada Business Corporations Act. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Bradshaw, M. T., C. Ma, B. Yost, and Y. Zou. 2025. Generative AI Use by Capital Market Information Intermediaries: Evidence from Seeking Alpha. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Breuer, M., W. Cai, A. Le, and F. Vetter. 2024. Minority Representation at Work. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Cai, W., Y. Chen, S. Rajgopal, and L. Azinovic-Yang. 2024. Diversity targets. *Review of Accounting Studies* 29 (3): 2157–2208.

Celiktutan, B., R. Cadario, and C. K. Morewedge. 2024. People see more of their biases in algorithms. *Proceedings of the National Academy of Sciences* 121 (16): e2317602121.

Chang, A., X. Dong, X. Martin, and C. Zhou. 2025. AI (ChatGPT) Democratization and Trading Inequality. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Chen, C.-W., and L. Y. Li. 2023. Is hiring fast a good sign? The informativeness of job vacancy duration for future firm profitability. *Review of Accounting Studies* 28 (3): 1316–1353.

Chang, E. H., and E. L. Kirgios. 2024. Demographic "Stickiness": The Demographic Identity of Departing Group Members Influences Who Is Chosen to Replace Them. *Management Science* 70 (7): 4236–4259.

Chen, Z. 2023. Ethics and discrimination in artificial intelligence-enabled recruitment practices. *Humanities and Social Sciences Communications* 10 (1): 567.

Choi, J. H., J. Pacelli, K. M. Rennekamp, and S. Tomar. 2023. Do jobseekers value diversity information? Evidence from a field experiment and human capital disclosures. *Journal of Accounting Research* 61 (3): 695–735.

Choi, J. H., and C. Xie. 2025. Human + AI in Accounting: Early Evidence from the Field. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Clement, M.B. 1999. Analyst forecast accuracy: Do ability, resources, and portfolio complexity matter? *Journal of Accounting and Economics* 27(3): 285-303.

Dastin, J. 2018. Insight- Amazon scraps secret AI recruiting tool that showed bias against women. Reuters, October 10.

DeAngelo, L. E. 1981. Auditor size and audit quality. *Journal of Accounting and Economics* 3 (3): 183–199.

DeFond, M., and J. Zhang. 2014. A review of archival auditing research. *Journal of Accounting and Economics* 58 (2–3): 275–326.

Dennison, K. 2024. Could lawsuits against AI lead to a shift in job searching? Forbes, March 21. https://www.forbes.com/sites/karadennison/2024/03/21/could-lawsuits-against-ai-lead-to-a-shift-in-job-searching/.

Dong, T. 2022. Gender Salary Gap in the Auditing Profession: Trend and Explanations.

European Accounting Review 33(2): 617–645.
Francis, J. R. 2004. What do we know about audit quality? *The British Accounting Review* 36 (4): 345–368.
Gassam Asare, J. 2025. What the Workday lawsuit reveals about AI bias – and how to prevent it. Forbes, June 23.
Gipper, B., S. Ross, and S. X. Shi. 2025. ESG Assurance in the United States. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.
Glum, J. 2023. Summer of strikes: why so many workers are walking off the job. Money, July 21. https://money.com/worker-strikes-summer-2023/
Godfrey, A. 2025. Understanding public perceptions towards automated decision-making in recruitment. *Revealing Reality*.
Goodman, R. 2018. Why Amazon's automated hiring tool discriminated against women. ACLA News & Commentary, October 12.
Hickock, M. 2022. NYC Bias Audit Law: Clock ticking for Employers and HR Talent Technology Vendors. Credo AI. August 15.
Hilliard, A., A. Gulley, A. Koshiyama, and E. Kazim. 2024. Bias audit laws: how effective are they at preventing bias in automated employment decision tools? International Review of Law, Computers & Technology, 1-17.
Jaser, Z., D. Petrakaki, R. Starr, and E. Oyarbide-Magaña. 2021. Automated job interviews and the implications for young jobseekers.
Khavis, J., A. Sheneman, and B. Szerwo. 2025. Does gender composition of audit teams matter? An examination of audit quality and audit cost. *Review of Accounting Studies*, forthcoming.
Kim, P. T. 2016. Data-Driven Discrimination at Work. *William & Mary Law Review* 58: 857.
Kim, P. T. 2017. Auditing Algorithms for Discrimination. *University of Pennsylvania Law Review Online* 166: 189.
Knechel, W.R., S. Maex, and H.J. Park. 2025. The role of auditor reputation in an emerging audit marketplace: Evidence from Decentralized Finance (DeFi). *Management Science*, forthcoming.
Kroll, J., J. Huey, S. Barocas, E. Felten, J. Reidenberg, D. Robinson, and H. Yu. 2017. Accountable Algorithms. *University of Pennsylvania Law Review* 165 (3): 633.
Lam, K., B. Lange, B. Blili-Hamelin, J. Davidovic, S. Brown, and A. Hasan. 2024. A Framework for Assurance Audits of Algorithmic Systems. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 1078–1092.
Larcker, D. F., C. McClure, S. X. Shi, and E. M. Watts. 2025. The Limited Corporate Response to DEI Controversies. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.
Lawrence, A., M. Minutti-Meza, and P. Zhang. 2011. Can Big 4 versus Non-Big 4 Differences in Audit-Quality Proxies Be Attributed to Client Characteristics? *The Accounting Review* 86 (1): 259–286.
Leicht-Deobald, U., T. Busch, C. Schank, A. Weibel, S. Schafheitle, I. Wildhaber, and G. Kasper. 2019. The Challenges of Algorithm-Based HR Decision-Making for Personal Integrity. *Journal of Business Ethics* 160 (2): 377–392.
Levy, B. 2025. Caution Ahead: Numerical Reasoning and Look-ahead Bias in AI Models. SSRN Scholarly Paper. Rochester, NY: Social Science Research Network.

Lins, K. V., L. Roth, H. Servaes, and A. Tamayo. 2024. Sexism, culture, and firm value: evidence from the Harvey Weinstein scandal and the# MeToo movement. *Journal of Accounting Research* 62 (5): 1989–2035.

Liu, T., C. A. Makridis, P. Ouimet, and E. Simintzi. 2023. The Distribution of Nonwage Benefits: Maternity Benefits and Gender Diversity. *The Review of Financial Studies* 36 (1): 194–234

Lockhart, J. W., M. M. King, and C. Munsch. 2023. Name-based demographic inference and the unequal distribution of misrecognition. *Nature Human Behaviour* 7 (7): 1084–1095.

Loh, R.K., and G.M. Mian. 2006. Do accurate earnings forecasts facilitate superior investment recommendations? *Journal of Financial Economics* 80(2): 455-483.

Matsa, D. A., and A. R. Miller. 2013. A Female Style in Corporate Leadership? Evidence from Quotas. *American Economic Journal: Applied Economics* 5 (3): 136–169.

McColl, R., and M. Michelotti. 2019. Sorry, could you repeat the question? Exploring video-interview recruitment practice in HRM. *Human Resource Management Journal* 29 (4): 637–656.

Mikhail, M.B., B.R. Walther, and R.H. Willis. 1999. Does forecast accuracy matter to security analysts? *The Accounting Review* 74(2): 185-200.

Milkman, R., and J. van der Naald. 2023. The State of the Unions 2023: A Profile of Organized Labor in New York City, New York State, and the United States. *Publications and Research*.

Moschella, D. 2022. *AI Bias Is Correctable. Human Bias? Not So Much*.

Munoko, I., H. L. Brown-Liburd, and M. Vasarhelyi. 2020. The Ethical Implications of Using Artificial Intelligence in Auditing. *Journal of Business Ethics* 167 (2): 209–234.

Oliveira, M., and S. Zhang. 2022. The trends and determinants of board gender and age diversities. *Finance Research Letters* 46: 102798.

Peng, A., B. Nushi, E. Kiciman, K. Inkpen, and E. Kamar. 2022. Investigations of Performance and Bias in Human-AI Teamwork in Hiring. *Proceedings of the AAAI Conference on Artificial Intelligence* 36 (11): 12089–12097.

Prince, A. E. R., and D. Schwarcz. 2019. Proxy Discrimination in the Age of Artificial Intelligence and Big Data. *Iowa Law Review* 105: 1257.

Raghavan, M., S. Barocas, J. Kleinberg, and K. Levy. 2020. Mitigating bias in algorithmic hiring: evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 469–481. FAT* ’20. New York, NY, USA: Association for Computing Machinery.

Raji, I. D., and J. 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. AIES ’19. New York, NY, USA: Association for Computing Machinery.

Schoenfeld, J. 2024. Cyber risk and voluntary Service Organization Control (SOC) audits. *Review of Accounting Studies* 29: 580-620.

Skitka, L. J., K.L. Mosier, and M. Burdick.1999. Does automation bias decision-making? *International Journal of Human-Computer Studies*, 51(5), 991-1006.

Suresh, H., and J. V. Guttag. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In *Equity and Access in Algorithms, Mechanisms, and Optimization*, 1–9.

Sutherland, A. G., M. Uckert, and F. W. Vetter. 2024. Occupational licensing and minority participation in professional labor markets. *Journal of Accounting Research* 62 (2): 453–503.

Van den Broek, E., A. Sergeeva, and M. Huysman. 2021. When the Machine Meets the Expert: An Ethnography of Developing AI for Hiring. *MIS Quarterly* 45 (3).

Wade, H. 2022. NYC Rules. https://rules.cityofnewyork.us/rule/automated-employment-decision-tools-2/.

Watts, R. L., and J. L. Zimmerman. 1983. Agency Problems, Auditing, and the Theory of the Firm: Some Evidence. *The Journal of Law and Economics* 26 (3): 613–633.

Wiessner. 2023. Tutoring firm settles US agency's first bias lawsuit involving AI software. https://www.reuters.com/legal/tutoring-firm-settles-us-agencys-first-bias-lawsuit-involving-ai-software-2023-08-10/.

Wiessner, D. 2024. Workday accused of facilitating widespread bias in novel AI lawsuit. *Reuters*, February 21, sec. Transactional.

Wilson, K., M. Sim, A.-M. Gueorguieva, and A. Caliskan. 2025. No Thoughts Just AI: Biased LLM Hiring Recommendations Alter Human Decision Making and Limit Human Autonomy. Proceedings of the AAAI ACM Conference on AI, Ethics, and Society, 8(3), 2692–2704.

Wooldridge, J.M. 2001. Econometric analysis of cross section and panel data. The MIT Press.

Wright, L., R. Mika Muenster, B. Vecchione, T. Qu, P. Cai, A. Smith, Comm 2450 Student Investigators, J. Metcalf, and J.N. Matias. 2024. Null Compliance: NYC Local Law 144 and the challenges of algorithm accountability. FAccT '24: Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency. 1701-1713.

Appendix Table A: Variable Definition

Variable Name	Variable Description
MaleProp	The proportion of male employees hired in the office during the quarter
WhiteProp	The proportion of white employees hired in the office during the quarter
WhiteMaleProp	The proportion of white male employees hired in the office during the quarter
NonWhiteMaleProp	The proportion of non-white male employees hired in the office during the quarter
LogEmpNum	The natural logarithm of the number of employees in the office during the quarter
Treat	Indicator variable equal to 1 if the firm's office is in New York City, and 0 otherwise
Post	Indicator variable equal to 1 for quarters on or after 2023 Q3 when the bias-audit law became effective, and 0 for quarters before
PropAgeUnder25	The proportion of employees under 25 years old hired in the office during the quarter
PropAge25To34	The proportion of employees aged 25-34 hired in the office during the quarter
PropAge35To44	The proportion of employees aged 35-44 hired in the office during the quarter
PropAge45Plus	The proportion of employees over 45 years old hired in the office during the quarter
Report	Indicator variable equal to 1 if a firm has issued a bias audit report, either *indirectly* (by adopting any of the 11 verified audited algorithms identified in the proprietary technology dataset, where the algorithm vendor's audit is assumed to extend to adopting firms) or *directly* (by publishing a bias audit report for an unspecified or uncovered algorithm), and 0 otherwise
NonReport	Indicator variable equal to 1 for a firm not classified as a *Report* firm, and 0 otherwise
MoreBiasedV	Indicator variable equal to 1 for firms using more biased algorithms among the 11 verified audited algorithms, and 0 otherwise. We classify an algorithm as more biased if its adopting firms exhibit a higher mean male hiring rate during the pre-Bias Audit Law period
LessBiasedV	Indicator variable equal to 1 for firms using less biased algorithms among the 11 verified audited algorithms, and 0 otherwise. We classify an algorithm as less biased if its adopting firms exhibit a lower mean male hiring rate during the pre-Bias Audit Law period
NY office	Indicator variable equal to 1 if the firm's office is in New York City, and 0 otherwise (equivalent to *Treat*)
Non-NY office	Indicator variable equal to 1 if the firm's office is not located in New York City, and 0 otherwise
LowHR	Indicator variable equal to 1 if the office's pre-law HR staff share is below the sample median, and 0 otherwise
HighHR	Indicator variable equal to 1 if the office's pre-law HR staff share is above the sample median
HRHireRate	The percentage of net HR staff hires, measured as the change in HR headcount from year $t-1$ to year t scaled by HR headcount in year $t-1$
LogDuration	The natural logarithm of the average number of days to fill a job vacancy in the office during the quarter
HireRate	The percentage of net hires in the office during the quarter, calculated as the change in the number of employees divided by prior-period number of employees

Variable Name	Variable Description
HighMaleFirm	Indicator variable equal to 1 if a firm's pre-law male hire share is above the sample median
LowMaleFirm	Indicator variable equal to 1 if a firm's pre-law male hire share is below the sample median
HighMaleInd	Indicator variable equal to 1 if a firm is in a male-dominated industry, defined as industries where the female employment share is below the overall average across all industries
LowMaleInd	Indicator variable equal to 1 if a firm is in a female-dominated industry, defined as industries where the female employment share is above the overall average across all industries
Error	The absolute difference between the analyst's quarterly EPS forecast and the quarterly EPS realized, scaled by the stock price
Overall Rating	The average overall rating for the office during the quarter, based on Glassdoor reviews from new hires (employees with less than one year of tenure), rated 1 to 5, with 5 being the best
Work-life balance Rating	The average work-life balance rating for the office during the quarter, based on Glassdoor reviews from new hires (employees with less than one year of tenure), rated 1 to 5, with 5 being the best
Senior Management Rating	The average senior management rating for the office during the quarter, based on Glassdoor reviews from new hires (employees with less than one year of tenure), rated 1 to 5, with 5 being the best
Diversity and Inclusion Rating	The average diversity and inclusion rating for the office during the quarter, based on Glassdoor reviews from new hires (employees with less than one year of tenure), rated 1 to 5, with 5 being the best
Career Opportunities Rating	The average career opportunities rating for the office during the quarter, based on Glassdoor reviews from new hires (employees with less than one year of tenure), rated 1 to 5, with 5 being the best
Culture and Values Rating	The average culture and values rating for the office during the quarter, based on Glassdoor reviews from new hires (employees with less than one year of tenure), rated 1 to 5, with 5 being the best
BlackProp	The proportion of Black employees hired in the office during the quarter
HispanicProp	The proportion of Hispanic employees hired in the office during the quarter
AsianProp	The proportion of Asian employees hired in the office during the quarter
Broker Size	The natural logarithm of the number of analysts employed by the brokerage house in the quarter
Analyst Following	The natural logarithm of the number of unique analysts issuing earnings forecasts for the firm in the quarter
Size	The natural logarithm of the firm's market value of equity in the quarter
Horizon	The natural logarithm of the number of days between the analyst forecast date and the corresponding earnings announcement date
BM	Book value of equity divided by market value of equity for the firm in the quarter
Lev	Total debt divided by total assets for the firm in the quarter
R&D	R&D expenditure scaled by total assets for the firm in the quarter
Earn Vol	The rolling standard deviation of quarterly earnings (income before extraordinary items scaled by lagged total assets) over the past 8 quarters, measured at the firm-quarter level

Variable Name	Variable Description
ROA	Income before extraordinary items divided by the total assets at the firm-quarter level
Inst Pct	Fraction of shares held by institutional investors.
Cash Flow	Cash flow from operations divided by the total assets

Figure 1: Timeline of the regulation

This figure presents the timeline of the New York City Bias Audit Law. The figure is sourced from https://www.nycbiasaudit.com/.

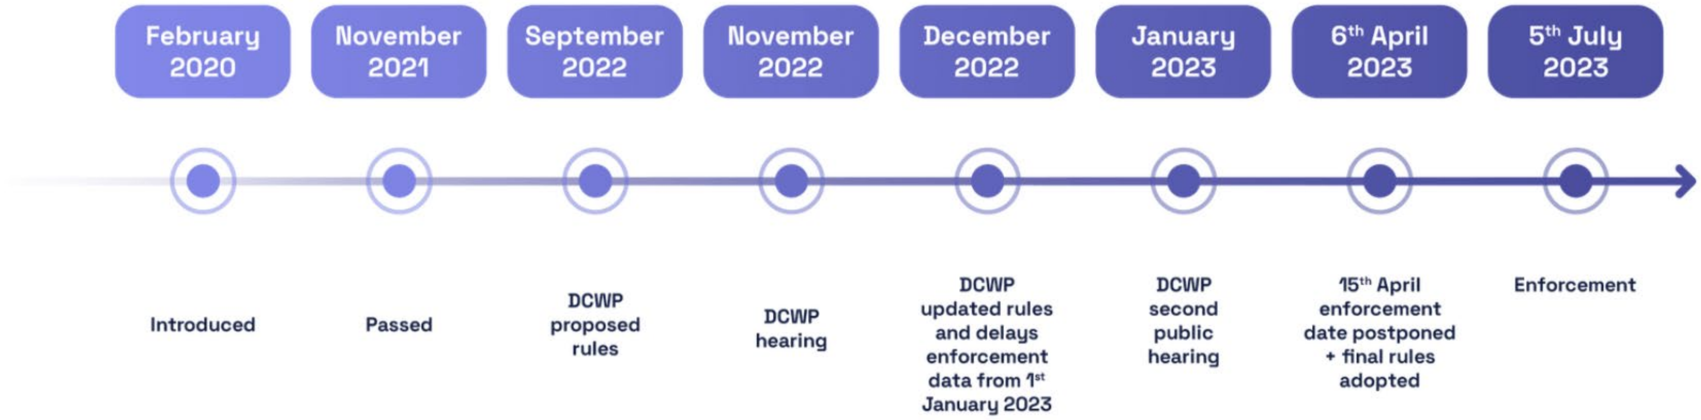

Figure 2: Research Design

This figure presents our research design. Panel A presents our sample window, starting from 2022 Q3 and ending with 2024 Q3. We use 2023 Q3, when the Bias Audit Law was enforced, as the event quarter (Event T= 0). Panel B presents a graphical summary of the identification strategy. Suppose Office 1 is located in New York City, while Office 2, from the same company, is located outside of New York City. Then Office 1 is in the treatment group while Office 2 is in the control group in our empirical design.

Panel A: Sample window

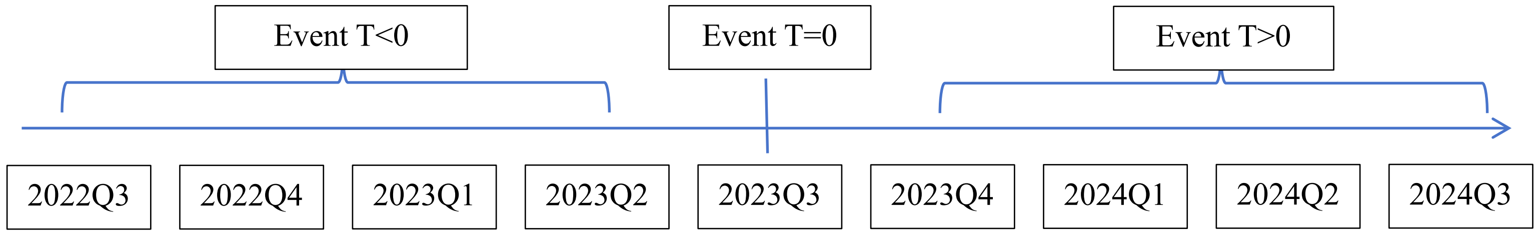

Panel B: Treatment and control group

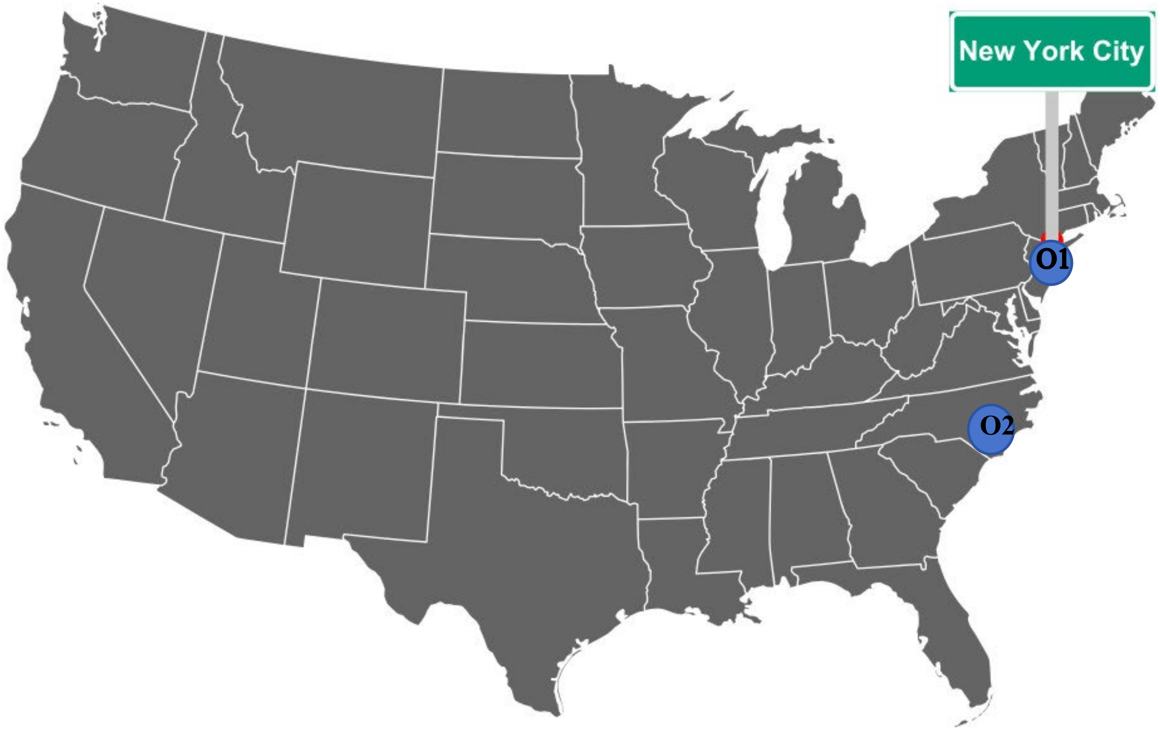

O1 (Treatment group): the office located in New York

O2 (Control group): the office of the same company located outside New York

Figure 3: Bias Audit Law and Male Proportion of Employment Hiring in Event-Time

This figure displays the dynamic effect of the Bias Audit Law on the male hiring proportion. The unit of observation is at the office-quarter level. The dependent variable is *MaleProp*, the proportion of male employees hired. The model replaces the "Post" indicator with a set of event-time dummies, each measured relative to the benchmark quarter t-1 (2023 Q2). Solid points plot the estimated coefficients, and vertical lines indicate 90 percent confidence intervals.

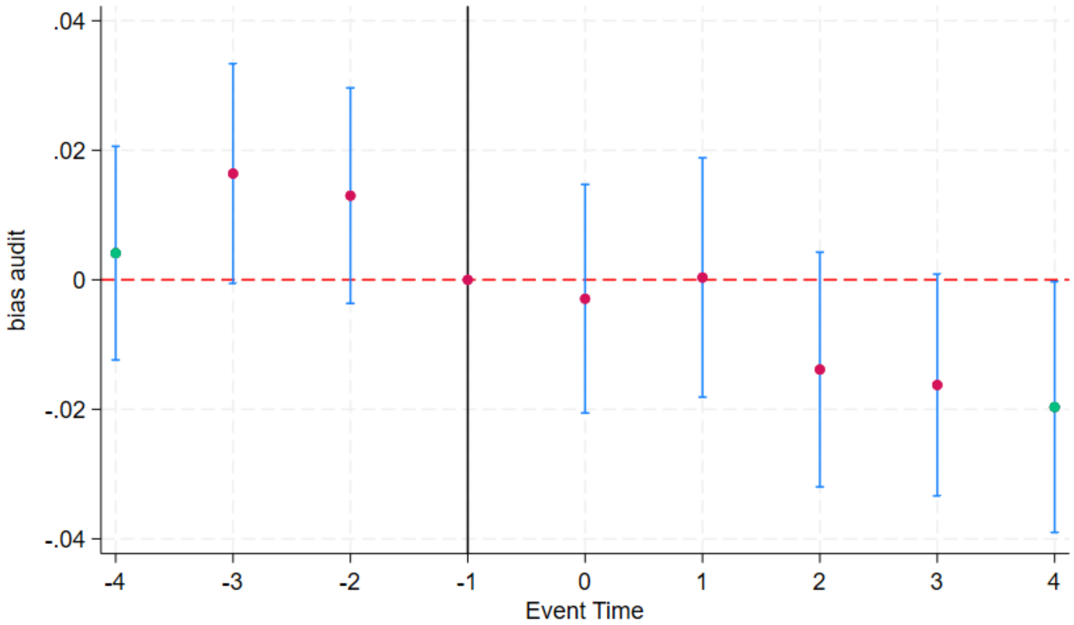

Figure 4: Bias Audit Law and HR Staff Net Growth Rate in Event-Time

This figure displays the dynamic effect of the Bias Audit Law on the HR staff net growth rate. The unit of observation is at the office-quarter level. The dependent variable *HRHireRate* is the net growth rate of HR staff, measured as the change in HR headcount from year t-1 to year t scaled by HR headcount in year t-1 The model replaces the "Post" indicator with a set of event-time dummies, each measured relative to the benchmark quarter t-1 (2023 Q2). Solid points plot the estimated coefficients, and vertical lines indicate 90 percent confidence intervals.

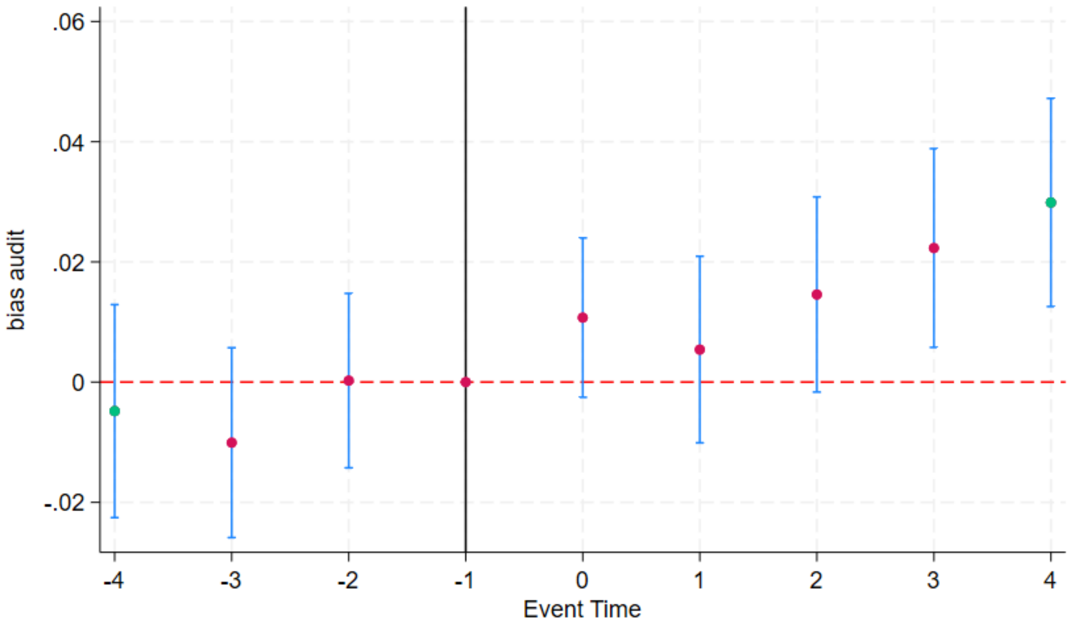

Figure 5: Bias Audit Law and Job Vacancy Duration in Event-Time

This figure displays the dynamic effect of the Bias Audit Law on the job vacancy duration. The unit of observation is at the office-quarter level. The dependent variable is *LogDuration*, the natural logarithm of the average number of days to fill a job vacancy in the office during the quarter. The model replaces the "Post" indicator with a set of event-time dummies, each measured relative to the benchmark quarter t-1 (2023 Q2). Solid points plot the estimated coefficients, and vertical lines indicate 90 percent confidence intervals.

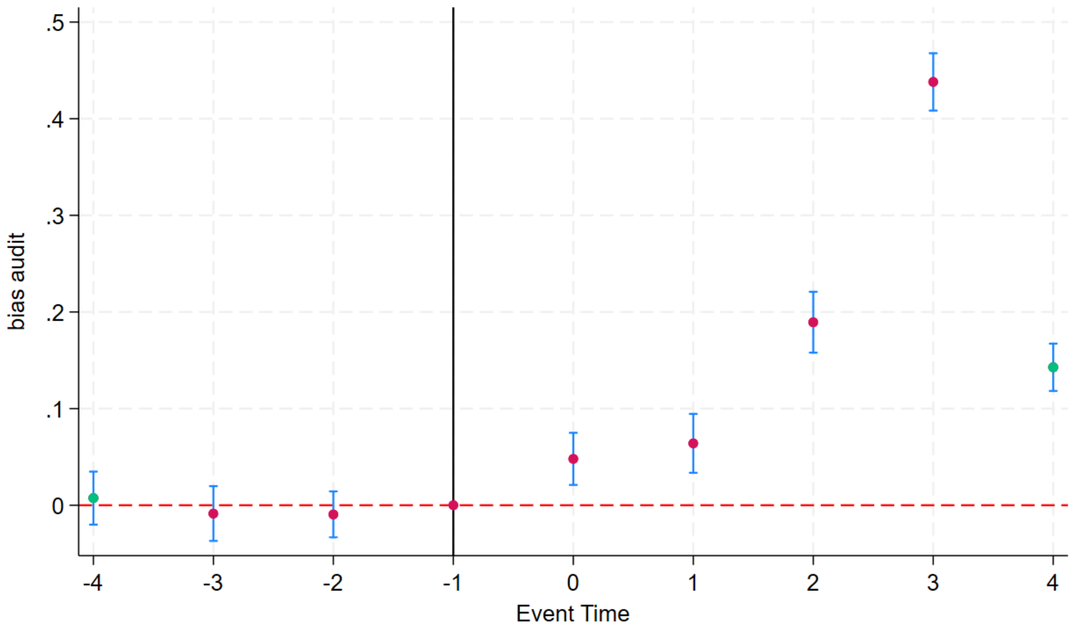

Table 1: Summary statistics

This table presents the summary statistics. Panel A reports the descriptive statistics for our main sample. Panel B reports the control-treatment difference of main dependent variables in the period prior to and after the enforcement of Bias Audit Law. Panels C and D further split New York City offices between audit reporters and non-reporters, respectively. All variables are defined in Appendix Table A.

Panel A: Descriptive Statistics for the Full Sample

	N	Mean	Median	SD	P25	P75
MaleProp	477,575	0.563	0.600	0.396	0.083	1.000
WhiteProp	477,575	0.699	0.857	0.362	0.500	1.000
WhiteMaleProp	477,575	0.396	0.333	0.389	0.000	0.667
NonWhiteMaleProp	477,575	0.166	0.000	0.289	0.000	0.250
LogEmpNum	477,575	4.811	4.727	1.807	3.524	6.057
PropAgeUnder25	352,222	0.351	0.200	0.395	0.000	0.667
PropAge25To34	352,222	0.347	0.250	0.381	0.000	0.571
PropAge35To44	352,222	0.180	0.000	0.311	0.000	0.250
PropAge45Plus	352,222	0.122	0.000	0.269	0.000	0.071
Report	477,575	0.414	0.000	0.492	0.000	1.000
NonReport	477,575	0.586	1.000	0.492	0.000	1.000
HRHireRate	221,138	-0.024	0.000	0.262	-0.061	0.000
LogDuration	578,374	3.539	3.584	0.562	3.258	3.885
HireRate	325,519	0.081	0.000	0.561	-0.315	0.078
Error	17,738	0.055	0.004	0.301	0.001	0.011
Overall Rating	49,682	3.523	4.000	1.230	3.000	4.500
Work-life balance Rating	44,042	3.482	3.800	1.306	2.750	5.000
Senior Management Rating	43,659	3.322	3.500	1.360	2.000	4.500
Diversity and Inclusion Rating	43,899	3.853	4.000	1.225	3.000	5.000
Career Opportunities Rating	44,342	3.452	3.667	1.299	2.667	4.667
Culture and Values Rating	44,073	3.504	4.000	1.340	2.667	5.000

Panel B: Treatment-control Comparison in both Pre- and Post-Period

	Treat			Control			
	Pre	Post	Diff *I* (Post - Pre)	Pre	Post	Diff *II* (Post - Pre)	DID (Diff *I* – Diff *II*)
MaleProp	0.581	0.567	-0.014**	0.557	0.567	0.010***	-0.024***
WhiteProp	0.634	0.642	0.008	0.697	0.706	0.009***	-0.002
HRHireRate	-0.022	-0.028	-0.005	-0.022	-0.034	-0.013***	0.007
LogDuration	3.545	3.764	0.219***	3.528	3.533	0.005***	0.214***

Panel C: Reporter (NYC offices) vs. control offices in both Pre- and Post-Period

	Report			Control			
	Pre	Post	Diff *I* (Post - Pre)	Pre	Post	Diff *II* (Post - Pre)	DID (Diff *I* – Diff *II*)
MaleProp	0.590	0.568	-0.022**	0.557	0.567	0.010***	-0.032***
WhiteProp	0.630	0.646	0.016*	0.697	0.706	0.009***	0.006
HRHireRate	-0.013	-0.020	-0.007	-0.022	-0.034	-0.013***	0.005
LogDuration	3.547	3.757	0.210***	3.528	3.533	0.005***	0.205***

Panel D: NonReporter (NYC offices) vs. control offices in both Pre- and Post-Period

	NonReport			Control			
	Pre	Post	Diff *II* (Post-Pre)	Pre	Post	Diff *II* (Post - Pre)	DID (Diff *I* – Diff *II*)
MaleProp	0.575	0.566	-0.008	0.557	0.567	0.010***	-0.019**
WhiteProp	0.638	0.639	0.001	0.697	0.706	0.009***	-0.008
HRHireRate	-0.031	-0.035	-0.004	-0.022	-0.034	-0.013***	0.008
LogDuration	3.544	3.770	0.226***	3.528	3.533	0.005***	0.221***

Table 2: Effect of the Bias Audit Law on Hiring Diversity

This table examines the effect of the Bias Audit Law on hiring diversity, using office-quarter observations. In Panel A, columns (1)–(2) report results for the proportion of male hires (*MaleProp*), and columns (3)–(4) for the proportion of white hires (*WhiteProp*). Panel B shows the proportion of white male hires (*WhiteMaleProp*) in column (1) and non-white male hires (*NonWhiteMaleProp*) in column (2). The primary explanatory variable is the interaction term between *Treat* and *Post*, where *Treat* equals 1 for offices located in New York City and *Post* equals 1 for the quarter the law takes effect and thereafter. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Proportion of Male Hires and Proportion of White Hires

	(1)	(2)	(3)	(4)
DV =	*MaleProp*	*MaleProp*	*WhiteProp*	*WhiteProp*
Treat* × *Post	**-0.018*** **	**-0.018*** **	-0.001	-0.001
	(-3.397)	**(-3.418)**	(-0.133)	(-0.145)
LogEmpNum		-0.005***		-0.003**
		(-3.523)		(-2.075)
Office FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes
R^2	0.371	0.371	0.374	0.374
N. of Obs.	477,575	477,575	477,575	477,575

Panel B: Proportion of White Male Hires and Proportion of Non-White Male Hires

	(1)	(2)
DV =	*WhiteMaleProp*	*NonWhiteMaleProp*
Treat* × *Post	**-0.010***	**-0.008***
	(-1.846)	**(-1.922)**
LogEmpNum	-0.005***	-0.001
	(-3.131)	(-0.506)
Office FE	Yes	Yes
Year-Quarter FE	Yes	Yes
R^2	0.349	0.336
N. of Obs.	477,575	477,575

Table 3: Placebo Test – Effect of the Bias Audit Law on Age Group Hiring

This table presents the results of a placebo test examining the effect of the Bias Audit Law on age-group hiring shares, using office-quarter observations. Columns (1)-(5) use, respectively, the proportion of hires aged under 25 (*PropAgeUnder25*), 25–34 (*PropAge25To34*), 35–44 (*PropAge35To44*), and 45 and over (*PropAge45Plus*) as the dependent variables. The primary explanatory variable is the interaction term between *Treat* and *Post*, where *Treat* equals 1 for offices located in New York City and *Post* equals 1 for the quarter the law takes effect and thereafter. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

	(1)	(2)	(3)	(4)
DV =	*PropAgeUnder25*	*PropAge25To34*	*PropAge35To44*	*PropAge45 Plus*
Treat* × *Post	-0.003	0.008	0.000	-0.005
	(-0.619)	(1.406)	(0.012)	(-1.258)
LogEmpNum	0.015***	-0.001	-0.007***	-0.006***
	(8.482)	(-0.680)	(-4.690)	(-4.947)
Office FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes
R^2	0.447	0.299	0.293	0.318
N. of Obs.	352,222	352,222	352,222	352,222

Table 4: Heterogeneous Effects between Audit-Reporters and Non-Reporters

This table examines the differential effect of the Bias Audit Law on the proportion of male hires between firms who conduct and do not conduct a bias audit, using office-quarter observations. The dependent variable is the proportion of male hires (*MaleProp*). *Treat* equals 1 for offices located in New York City, and *Post* equals 1 for the quarter the law takes effect and thereafter. Panel A compares the treatment effects between audit-reporters and non-reporters. *Report* equals one if a firm has issued bias audit reports, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. *NonReport* equals 1 for a firm subject to the Bias Audit Law but not classified as a *Report* firm. Panel B further splits reporters by the degree of pre-law gender bias in their AI hiring system to examine whether the audit effect is stronger for algorithms that are more biased ex-ante. *More(Less)BiasedV* equals 1 for algorithms whose adopting firms exhibited higher (lower) mean male hiring rates before the law. Panel C examines whether the auditing effect spills over to reporting firms' offices outside New York City. *NY office* equals one for firm's office located in New York City (equivalent to *Treat*). *Non-NY office* equals one for offices located outsider New York City. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Audit Reporters versus Non-Reporters

DV =	(1) *MaleProp*
Treat **×** ***Post*** **×** ***Report*** **(i)**	**-0.029***** **(-3.741)**
Treat **×** ***Post*** **×** ***NonReport*** **(ii)**	-0.010 (-1.419)
LogEmpNum	-0.005*** (-3.517)
F-statistics for (i) = (ii)	***3.10****
Office FE	Yes
Quarter FE	Yes
R^2	0.371
N. of Obs.	477,575

Panel B: Potentially More Biased Systems vs. Potentially Less Biased Systems

DV =	(1) *MaleProp*
Treat **×** ***Post*** **×** ***Report*** **×** ***MoreBiasedV*** **(i)**	**-0.049***** **(-3.542)**
Treat **×** ***Post*** **×** ***Report*** **×** ***LessBiasedV*** **(ii)**	**-0.018**** **(-1.963)**
Treat **×** ***Post*** **×** ***NonReport***	-0.010 (-1.418)
LogEmpNum	-0.005*** (-3.520)
F-statistics for (i) = (ii)	***3.47****
Office FE	Yes
Year-Quarter FE	Yes
R^2	0.371
N. of Obs.	477,575

Panel C: Spillover Effects to Non NYC Offices

DV =	(1) *MaleProp*
***NY office* × *Post* × *Report* (i)**	**-0.030***** **(-3.924)**
***Non-NY office* × *Post* × *Report* (ii)**	**-0.004*** **(-1.790)**
***NY office* × *Post* × *NonReport* (iii)**	-0.012 (-1.629)
LogEmpNum	-0.005*** (-3.484)
F- statistics for (i) = (ii)	***11.53****
F- statistics for (i) = (iii)	***3.10**
Office FE	Yes
Year-Quarter FE	Yes
R^2	0.371
N. of Obs.	477,575

Table 5: Heterogeneous Effects by Pre-Law HR Staff Share

This table examines how human oversight in AI hiring shapes the effect of the Bias Audit Law, using office-quarter observations. We use the pre-law HR staff share to proxy for the extent of human involvement in hiring, where firms with a lower HR staff share may rely more on AI-assisted hiring. *Treat* equals 1 for offices located in New York City, and *Post* equals 1 for the quarter the law takes effect and thereafter. The dependent variable in Panel A and B is the proportion of male hires (*MaleProp*). Panel A reports heterogeneous effects by pre-law HR staff share. *Low(High)HR* equals 1 for offices whose pre-law HR staff share is below (above) the sample median. Panel B reports heterogeneous effects along two dimensions: pre-law HR staff share and audit-reporter status. *Report* equals 1 if a firm has issued a bias audit report, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. *NonReport* equals one for a firm not classified as a *Report* firm. Panel C reports the percentage of net HR staff hires following the adoption of the bias audit law. The dependent variable is *HRHireRate*, the percentage of net HR hires, measured as the change in HR headcount from year t−1 to year t scaled by HR headcount in year t−1. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Heterogeneous Effects by HR Staff Share

	(1)
DV =	*MaleProp*
***Treat* × *Post* × *LowHR* (i)**	**-0.038*** **
	(-3.502)
***Treat* × *Post* × *HighHR* (ii)**	-0.009
	(-1.515)
LogEmpNum	-0.005***
	(-3.189)
F-statistics for (i) = (ii)	***5.89*** **
Office FE	Yes
Year-Quarter FE	Yes
R^2	0.367
N. of Obs.	465,903

Panel B: Heterogeneous Effects by HR Staff Share and Reporter Status

	(1)
DV =	*MaleProp*
***Treat* × *Post* × *Report* × *LowHR* (i)**	**-0.075*** **
	(-4.074)
***Treat* × *Post* × *Report* × *HighHR* (ii)**	-0.012
	(-1.518)
***Treat* × *Post* × *NonReport* × *LowHR* (iii)**	-0.019
	(-1.390)
***Treat* × *Post* × *NonReport* × *HighHR* (iv)**	-0.006
	(-0.688)
LogEmpNum	-0.005***
	(-3.178)
F-statistics for (i) = (ii)	***9.98*** **
F-statistics for (i) = (iii)	***6.07*** **
F-statistics for (ii) = (iv)	*0.33*
F-statistics for (iii) = (iv)	*0.72*
Office FE	Yes
Year-Quarter FE	Yes
R^2	0.367
N. of Obs.	465,903

Panel C: Net HR Staff Hires

DV =	(1) *HRHireRate*	(2) *HRHireRate*
Treat* × *Post	**0.020***** **(3.542)**	
***Treat* × *Post* × *Report* (i)**		**0.025***** **(3.179)**
***Treat* × *Post* × *NonReport* (ii)**		**0.014*** **(1.875)**
LogEmpNum	0.026*** (13.794)	0.026*** (13.789)
F-statistics for (i) = (ii)		*1.09*
Office FE	Yes	Yes
Year-Quarter FE	Yes	Yes
R^2	0.279	0.279
N. of Obs.	221,138	221,138

Table 6: Effect of Bias Audit Law on Job Vacancy Duration

This table examines the effect of the Bias Audit Law on job vacancy duration, using office-quarter observations. Column (1) reports OLS estimates with *LogDuration*, the natural logarithm of the average days to fill a vacancy in a quarter, as the dependent variable. Column (2) reports Poisson estimates with *Duration*, the average days to fill a vacancy in a quarter, as the dependent variable. *Treat* equals 1 for offices located in New York City, and *Post* equals 1 for the quarter the law takes effect and thereafter. Panel A reports the results on the job vacancy duration. Panel B further compares the effects between audit reporters and non-reporters. *Report* equals 1 if a firm has issued a bias audit report, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. *NonReport* equals one for a firm not classified as a *Report* firm. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Job Vacancy Duration

	(1)	(2)
	OLS Regression	Poisson Regression
DV =	*LogDuration*	*Duration*
Treat* × *Post	**0.189*** **	**0.200*** **
	(18.535)	**(21.134)**
LogEmpNum	0.062***	-0.038***
	(25.686)	(-16.954)
Office FE	Yes	Yes
Year-Quarter FE	Yes	Yes
R^2 (Pseudo R^2)	0.433	0.294
N. of Obs.	578,374	578,374

Panel B: Job Vacancy Duration by Reporter Status

	(1)	(2)
	OLS Regression	Poisson Regression
DV =	*LogDuration*	*Duration*
***Treat* × *Post* × *Report* (i)**	**0.180*** **	**0.194*** **
	(12.160)	**(13.954)**
***Treat* × *Post* × *NonReport* (ii)**	**0.196*** **	**0.204*** **
	(14.188)	**(16.108)**
LogEmpNum	0.062***	-0.038***
	(25.686)	(-16.953)
F-statistics for (i) = (ii)	*0.64*	*0.26*
Office FE	Yes	Yes
Year-Quarter FE	Yes	Yes
R^2 (Pseudo R^2)	0.433	0.294
N. of Obs.	578,374	578,374

Table 7: Effect of Bias Audit Law on Total Hires

This table examines the effect of the Bias Audit Law on offices' overall hiring rate, using office-quarter observations. The dependent variable is *HireRate*, the percentage change in the number of employees relative to the prior period. *Treat* equals 1 for offices located in New York City, and *Post* equals 1 for the quarter the law takes effect and thereafter. Panel A reports the baseline treatment effect on hiring rate. Panel B further compares the effects between audit reporters and non-reporters. *Report* equals 1 if a firm has issued a bias audit report, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. *NonReport* equals 1 for firms not classified as reporters. Panel C examines whether firms shift hiring activity to their neighboring offices outside New York City. *NeighborOffice* equals 1 for other offices within the New York MSA but outside New York State. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Overall Hiring Rate

DV =	(1) *HireRate*
Treat* × *Post	-0.009
	(-1.015)
LogEmpNum	1.036***
	(274.903)
Office FE	Yes
Year-Quarter FE	Yes
R^2	0.497
N. of Obs.	325,519

Panel B: Hiring Rate by Reporter Status

DV =	(1) *HireRate*
Treat* × *Post* × *Report	-0.020
	(-1.498)
Treat* × *Post* × *NonReport	0.000
	(0.008)
LogEmpNum	1.036***
	(274.888)
Office FE	Yes
Year-Quarter FE	Yes
R^2	0.497
N. of Obs.	325,519

Panel C: Hiring Displacement to Neighboring Offices

DV =	(1) *HireRate*
Treat* × *Post	-0.009
	(-0.998)
NeighborOffice* × *Post	0.008
	(0.581)
LogEmpNum	1.036***
	(274.900)
Office FE	Yes
Year-Quarter FE	Yes
R^2	0.497
N. of Obs.	325,519

Table 8: Heterogeneous Effects by Pre-Law Male Dominance

This table examines the differential effect of the Bias Audit Law on the proportion of male hires, net hires of HR staff, and job vacancy duration across firms operating in more versus less male-dominated environments, using office-quarter observations. We proxy for male dominance using two measures: The first is at the firm level: *High(Low)MaleFirm* equals 1 if a firm's pre-law male hiring share is above (below) the sample median. The second is at the industry level: *High(Low)MaleInd* equals 1 if a firm operates in an industry where the female employment share is below (above) the overall average across all industries. Columns (1) and (2) use *MaleProp*, the proportion of male hires, as the dependent variable. Columns (3) and (4) use *HRHireRate*, the percentage of net HR hires, as the dependent variable. Columns (5) and (6) use *LogDuration*, the natural logarithm of the average days to fill a vacancy in a quarter, as the dependent variable. *Treat* equals 1 for offices located in New York City, and *Post* equals 1 for the quarter the law takes effect and thereafter. Panel A compares the treatment effects between firms with higher versus lower degrees of male dominance. Panel B further examines whether these heterogeneous effects differ between audit reporters and non-reporters. *Report* equals 1 if a firm has issued a bias audit report, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. *NonReport* equals one for a firm not classified as a *Report* firm. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Male-dominated and Non-male-dominated Firms (Industries)

	(1)	(2)	(3)	(4)	(5)	(6)
DV =	*MaleProp*	*MaleProp*	*HRHireRate*	*HRHireRate*	*LogDuration*	*LogDuration*
***Treat* × *Post* × *HighMaleFirm* (i)**	**-0.031*** **		**0.020*** **		**0.222*** **	
	(-4.130)		**(3.021)**		**(17.546)**	
***Treat* × *Post* × *LowMaleFirm* (ii)**	-0.004		**0.019***		**0.122*** **	
	(-0.497)		**(1.914)**		**(7.202)**	
***Treat* × *Post* × *HighMaleInd* (iii)**		**-0.031*** **		**0.028*** **		**0.224*** **
		(-4.861)		**(3.688)**		**(15.323)**
***Treat* × *Post* × *LowMaleInd* (iv)**		0.009		0.011		**0.152*** **
		(0.945)		(1.366)		**(10.878)**
LogEmpNum	-0.005***	-0.005***	0.026***	0.026***	0.062***	0.062***
	(-3.520)	(-3.497)	(13.825)	(13.794)	(25.674)	(25.672)
F-statistics for (i) = (ii)	***6.98**** **		*0.01*		***22.51**** **	
F-statistics for (iii) = (iv)		***12.10**** **		***2.77****		***12.89**** **
Office FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.371	0.371	0.279	0.279	0.433	0.433
N. of Obs.	477,575	477,174	221,051	221,138	577,151	578,374

Panel B: Heterogeneous Effects by Male Dominance and Reporter Status

DV =	(1) *MaleProp*	(2) *MaleProp*	(3) *HRHireRate*	(4) *HRHireRate*	(5) *LogDuration*	(6) *LogDuration*
***Treat* × *Post* × *Report* × *HighMaleFirm* (i)**	**-0.039*****		**0.022****		**0.199*****	
	(-3.811)		**(2.330)**		**(11.044)**	
***Treat* × *Post* × *Report* × *LowMaleFirm* (ii)**	-0.013		**0.033****		**0.125*****	
	(-1.215)		**(2.277)**		**(5.016)**	
***Treat* × *Post* × *NonReport* × *HighMaleFirm* (iii)**	**-0.024****		**0.018****		**0.239*****	
	(-2.187)		**(1.972)**		**(13.726)**	
***Treat* × *Post* × *NonReport* × *LowMaleFirm* (iv)**	0.003		0.006		**0.120*****	
	(0.262)		(0.492)		**(5.314)**	
***Treat* × *Post* × *Report* × *HighMaleInd* (v)**		**-0.041*****		**0.032*****		**0.200*****
		(-4.615)		**(2.932)**		**(9.940)**
***Treat* × *Post* × *Report* × *LowMaleInd* (vi)**		0.003		0.016		**0.154*****
		(0.191)		(1.458)		**(7.160)**
***Treat* × *Post* × *NonReport* × *HighMaleInd* (vii)**		**-0.022****		**0.024****		**0.242*****
		(-2.527)		**(2.280)**		**(11.757)**
***Treat* × *Post* × *NonReport* × *LowMaleInd* (viii)**		0.013		0.006		**0.150*****
		(1.051)		(0.573)		**(8.281)**
LogEmpNum	-0.005***	-0.005***	0.026***	0.026***	0.062***	0.062***
	(-3.515)	(-3.491)	(13.822)	(13.790)	(25.673)	(25.669)
F-statistics for (i) = (ii)	***2.97****		*0.40*		***5.91*****	
F-statistics for (i) = (iii)	*1.16*		*0.10*		*2.49*	
F-statistics for (ii) = (iv)	*1.20*		*1.88*		*0.02*	
F-statistics for (iii) = (iv)	***3.30****		*0.53*		***17.27******	
F-statistics for (v) = (vi)		***6.64******		*0.99*		*2.42*
F-statistics for (v) = (vii)		*2.37*		*0.28*		*2.16*
F-statistics for (vi) = (viii)		*0.30*		*0.48*		*0.02*
F-statistics for (vii) = (viii)		***5.32*****		*1.51*		***11.36******
Office FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.371	0.371	0.279	0.279	0.433	0.433
N. of Obs.	477,575	477,575	221,051	221,138	577,151	578,374

Table 9: Effect of Bias Audit Law on Firm Performance: Evidence from Brokerage Firms

This table aims to examine whether improvements in hiring diversity driven by the Bias Audit Law translate into better firm performance, using a brokerage firm subsample. Panel A validates that the main hiring diversity findings hold in the brokerage subsample by replicating the effects on the proportion of male and white hires, HR staff growth rate, and job vacancy duration. Columns (1) estimates the effect on male hiring share (*MaleProp*), columns (2) on white hiring share (*WhiteProp*), column (3) on net growth rate of HR staff (*HRHireRate*), and (4) on job vacancy duration (*LogDuration*). Panels B and C examine whether the Bias Audit Law improves analyst forecast accuracy for analysts based in New York City offices. Panel B reports the summary statistics and Panel C reports the regression estimates. The unit for analysis is at the analyst-stock-quarter level. The dependent variable is *Error*, which is the absolute difference between the analyst's quarterly EPS forecast and realized EPS, scaled by the stock price. The primary explanatory variable is the interaction term between *Treat* and *Post*, where *Treat* equals 1 for analysts based in a New York City office of the brokerage firm and *Post* equals 1 for the month (quarter) the law takes effect and thereafter. *Report* equals 1 if a firm has issued a bias audit report, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. *NonReport* equals 1 for firms not classified as reporters. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Employment hiring

	(1)	(2)	(3)	(4)
DV =	*MaleProp*	*WhiteProp*	*HRHireRate*	*LogDuration*
Treat* × *Post	**-0.040****	-0.019	**0.039****	**0.169*****
	(-1.994)	(-1.042)	**(2.208)**	**(2.745)**
LogEmpNum	-0.013	0.008	0.028***	0.004
	(-1.388)	(0.916)	(3.071)	(0.276)
Office FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	No	Yes
Year-Month FE	No	No	Yes	No
R^2(Pseudo R^2)	0.280	0.370	0.267	0.296
N. of Obs.	15,137	15,137	7,616	13,270

Panel B: Summary Statistics

	N	Mean	Median	SD	P25	P75
Error	17,738	0.055	0.004	0.301	0.001	0.011
Broker Size	17,738	4.146	4.290	0.779	3.526	4.868
Analyst Following	17,738	10.509	10.000	6.211	6.000	14.00
Size	17,738	8.823	8.824	1.907	7.585	10.204
Horizon	17,738	3.546	3.850	0.978	2.890	4.431
BM	17,738	0.479	0.366	0.488	0.152	0.692
Lev	17,738	0.333	0.293	0.250	0.136	0.469
R&D	17,738	0.015	0.000	0.029	0.000	0.016
Earn Vol	17,738	0.020	0.010	0.031	0.005	0.022
ROA	17,738	0.005	0.014	0.050	0.002	0.030
Inst Pct	17,738	0.015	0.000	0.030	0.000	0.000
Cash Flow	17,738	0.021	0.026	0.104	0.000	0.069

Panel C: Analyst Forecast Error

DV =	(1) *Error*	(2) *Error*
Treat* × *Post	**-0.010****	
	(-2.784)	
Treat* × *Post* × *Report		**-0.018*****
		(-4.717)
Treat* × *Post* × *NonReport		-0.005
		(-1.372)
Broker Size	-0.006**	-0.004
	(-2.347)	(-1.585)
Analyst Following	-0.000	0.000
	(-0.068)	(0.059)
Size	-0.049***	-0.050***
	(-5.519)	(-5.576)
Horizon	0.001	0.001
	(1.514)	(1.430)
BM	0.057***	0.057***
	(3.923)	(3.911)
Lev	0.026	0.026
	(0.806)	(0.815)
R&D	0.097	0.094
	(0.408)	(0.393)
Earn Vol	0.365*	0.370*
	(1.815)	(1.830)
ROA	-0.971***	-0.973***
	(-15.969)	(-15.926)
Inst Pct	-0.282***	-0.282***
	(-6.643)	(-6.638)
Cash Flow	0.105***	0.104***
	(4.756)	(4.758)
F-statistics for (i) = (ii)		***491.68******
Stock FE	Yes	Yes
Year-Quarter FE	Yes	Yes
Analyst FE	Yes	Yes
R^2	0.856	0.856
N. of Obs.	17,738	17,738

Table 10: Effect of Bias Audit Law on Employee Satisfaction: Evidence from Glassdoor Reviews

This table examines whether the Bias Audit Law affects employee satisfaction measured using Glassdoor ratings from new hires (employees with less than one year of tenure), using office-quarter observations. Panel A reports the baseline treatment effects across six rating categories. Panel B further compares the effects between audit reporters and non-reporters to examine whether improvements in employee satisfaction are concentrated among firms that have audited their AI hiring systems. *Overall Rating* is the average employee overall rating for the office during the quarter. *Work-life balance* is the average work-life balance rating for the office during the quarter. *Senior Management* is the average senior management rating for the office during the quarter. *Diversity and Inclusion* is the average diversity and inclusion rating for the office during the quarter. *Career Opportunities* is the average career opportunities rating for the office during the quarter. *Culture and Values* is the average culture and values rating for the office during the quarter. *Treat* equals 1 for offices located in New York City, and *Post* equals 1 for the quarter the law takes effect and thereafter. *Report* equals 1 if a firm has issued a bias audit report, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. *NonReport* equals 1 for firms not classified as reporters. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Glassdoor Reviews

	(1)	(2)	(3)	(4)	(5)	(6)
DV=	*Overall Rating*	*Work-life balance*	*Senior Management*	*Diversity and Inclusion*	*Career Opportunities*	*Culture and Values*
Treat × Post	**0.081***	0.011	0.040	**0.089***	**0.084***	0.077
	(1.818)	(0.223)	(0.752)	**(1.905)**	**(1.788)**	(1.513)
NumReview	0.012***	0.013***	0.017***	0.015***	0.017***	0.014***
	(3.087)	(3.018)	(3.726)	(3.810)	(4.169)	(3.276)
Office FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.397	0.401	0.395	0.360	0.408	0.404
N. of Obs.	49,682	44,042	43,659	43,899	44,342	44,073

Panel B: Glassdoor Reviews vs. Compliers and Non-Compliers

DV=	(1) *Overall Rating*	(2) *Work-life balance*	(3) *Senior Management*	(4) *Diversity and Inclusion*	(5) *Career Opportunities*	(6) *Culture and Values*
Treat* × *Post* × *Report	**0.158*****	0.056	0.064	**0.118***	0.099	**0.144****
	(2.598)	(0.822)	(0.862)	**(1.834)**	(1.444)	**(2.051)**
Treat* × *Post* × *NonReport	0.003	-0.034	0.015	0.061	0.069	0.009
	(0.050)	(-0.507)	(0.212)	(0.922)	(1.118)	(0.135)
NumReview	0.012***	0.013***	0.017***	0.015***	0.017***	0.014***
	(3.097)	(3.022)	(3.728)	(3.812)	(4.170)	(3.281)
F-statistics for (i) = (ii)	***3.23****	*0.92*	*0.23*	*0.41*	*0.11*	*1.90*
Office FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.397	0.401	0.395	0.360	0.408	0.404
N. of Obs.	49,682	44,042	43,659	43,899	44,342	44,073

Online Appendix to

Auditing Effects on Employment Hiring: Evidence from the New York City Algorithmic Bias Audit Law

Online Appendix A: A Sample of Audit Report

This appendix presents a sample audit report issued by BABL for the bias audit results of Eightfold.

Bias Audit for New York City Local Law 144
Prepared by BABL AI Inc. | 03/21/2025
Letter from the Lead Auditor | Summary | Conclusions | Findings

babl

Letter from the Lead Auditor

From: **Shea Brown**
Lead Auditor
BABL AI Inc.
sheabrown@bablai.com

To: **Eightfold AI Inc.**
2625 Augustine Drive
Suite 601
Santa Clara, CA 95054

Re: **Audit Opinion on Eightfold AI Inc.'s Eightfold Matching Model**

03/21/2025

We have independently audited the bias testing assertions and related documentary evidence of Eightfold AI Inc. (the "Company") as of 03/21/2025, presented to BABL AI in relation to Company's Eightfold Matching Model in accordance with the criteria and audit methodology set forth in this report. The goals of this audit are to:

1. Determine whether the bias testing methodologies, controls, and procedures performed by Company satisfy the audit criteria (see Findings)
2. Obtain reasonable assurance as to whether the statements made by the Company, including the summary of bias testing results presented in this report, are free from material misstatement, whether due to fraud or error.

Note that the criteria presented in this report were constructed specifically to address the requirements of a "bias audit" outlined in NYC Local Law No. 144 of 2021. The model was audited as though it were an automated employment decision tool (AEDT) under NYC Local Law No. 144 of 2021, but we do not make any determination whether the model is, in fact, an AEDT under this law.

Company Responsibilities

It is the responsibility of Company representatives to ensure that bias testing and related procedures comply with the criteria outlined in this report. The Company representatives are responsible for ensuring that the documents submitted are fairly presented and free of misrepresentations, providing all resources and personnel needed to ensure an effective and efficient audit process, and providing access to evidential material as requested by the auditors.

BABL AI Responsibilities

It is the responsibility of the lead auditor to express an opinion on the Company's assertions related to the bias testing of the model. In light of the current absence of generally accepted standards for the auditing of algorithms and autonomous systems, our examination was conducted in accordance with the standards and normative references outlined in this report.

Those standards require that we plan and perform audit procedures to obtain reasonable assurance about whether the assertions referred to above 1) satisfy the audit criteria and 2) are free of material misstatement, whether due to error or fraud. Within the scope of our engagement, we performed amongst others the following procedures:

- Inspection of submitted documents and external documentation
- Interviewing Company employees to gain an understanding of the process for determining the disparate impact and risk assessment results
- Observation of selected analytical procedures used in Company's bias testing
- Inspection of the select samples of the bias testing data
- Inquiry of personnel responsible for governance and oversight of the bias testing and risk assessment

We believe that the procedures performed provide a reasonable basis for our opinion.

Independence

Our role as an independent auditor conforms to ForHumanity and Sarbanes-Oxley definitions of Independence. Fees associated with this contract are for the provision of the service to assess compliance. The payment of fees is unrelated to the decision rendered. Our decision is grounded solely in the criteria presented below.

Opinion

In our opinion, based on the procedures performed and the evidence received to obtain assurance, the bias testing and results presented by Company, as of 03/21/2025, is prepared, in all material respects, in accordance with the criteria outlined below.

Sincerely,

Shea Brown

Shea Brown
Lead Auditor, BABL AI Inc.

Non-intersectional, Gender, sorted by Scoring rate

	N applicants	Scoring rate	Impact ratio
Female	116	0.509	1.000
Male	739	0.491	0.966

Non-intersectional, Race/ethnicity, sorted by Scoring rate

	N applicants	Scoring rate	Impact ratio[7]
Black or African American	44	0.545	1.000
Hispanic or Latino	48	0.521	0.955
White	174	0.506	0.927
Asian	589	0.484	0.887

Intersectionals

<table>
<tr><th></th><th></th><th></th><th>N applicants</th><th>Scoring rate</th><th>Impact ratio[8]</th></tr>
<tr><td rowspan="2">Hispanic or Latino</td><td colspan="2">Male</td><td>43</td><td>0.512</td><td>0.853</td></tr>
<tr><td colspan="2">Female</td><td>5</td><td>0.600</td><td>1.000</td></tr>
<tr><td rowspan="10">Non-Hispanic or Latino</td><td rowspan="6">Male</td><td>White</td><td>151</td><td>0.503</td><td>0.839</td></tr>
<tr><td>Asian</td><td>506</td><td>0.482</td><td>0.804</td></tr>
<tr><td>Black or African American</td><td>39</td><td>0.538</td><td>0.897</td></tr>
<tr><td>Native American or Alaskan Native</td><td>0</td><td>0.000</td><td>N/A</td></tr>
<tr><td>Native Hawaiian or Pacific Islander</td><td>0</td><td>0.000</td><td>N/A</td></tr>
<tr><td>Two or more races</td><td>0</td><td>0.000</td><td>N/A</td></tr>
<tr><td rowspan="4">Female</td><td>Asian</td><td>83</td><td>0.494</td><td>0.823</td></tr>
<tr><td>White</td><td>23</td><td>0.522</td><td>0.870</td></tr>
<tr><td>Black or African American</td><td>5</td><td>0.600</td><td>1.000</td></tr>
<tr><td>Native American or Alaskan Native</td><td>0</td><td>0.000</td><td>N/A</td></tr>
</table>

Online Appendix B

This online appendix presents the results of additional analyses briefly discussed in the paper. Table A1 tests whether AI hiring systems are on average biased toward male hires prior to the Bias Audit Law, using office-quarter observations for the NYC office subsample. Table A2 focuses on the effect on male hires partitioned by job seniority. Table A3 on the effect on male hires partitioned by the public listing status of the employer, and Table A4 on heterogeneous effects on non-white subgroup hiring shares.

Job Seniority

We examine whether the impact of the Bias Audit Law varies by job seniority level. Senior positions often require specific skills and extensive experience, which lowers the pool of eligible candidates and can limit the applicability of standardized automated screening tools. These roles may also rely more heavily on networking and executive search processes that fall outside the scope of algorithmic decision tools. In contrast, junior positions typically involve more standardized hiring processes with clearly defined qualifications and much larger volumes of applications, making them more suitable for AI-based screening (Godfrey, 2025; Jaser et al., 2021). In addition, junior roles entail lower adjustment costs when firms respond to fairness mandates (Chang and Kirgios, 2024; Larcker et al., 2025). Therefore, we expect that the decline in male hiring is primarily concentrated in junior positions.

To test this prediction, we conduct subsample analyses based on seniority levels in Revelio Labs, where 1 = entry, 2 = junior, 3 = associate, 4 = manager, 5 = director, 6 = executive, and 7 = senior executive. Table A2 Panel A presents the results of estimating separate regressions for entry-level positions (seniority = 1 or 2) in column (1), Mid-level positions (seniority = 3, 4, or 5) in column (2), and senior executive positions (seniority = 6 or 7) in column (3). The coefficients on *Treat*×*Post* are significantly negative only in columns

(1) and (2). These results suggest that the reduction in male hiring is particularly pronounced in lower to mid-tier positions, likely because AI screening tools are more relevant for these positions or diversity adjustments are easier to implement among these roles.

This analysis also helps alleviate concerns that NYC firms respond more to national trends—such as corporate diversity fellowships, on-campus recruiting expansions, unionization drives, and skills-based hiring initiatives—due to the city's unique culture. Such initiatives disproportionately target junior-level or lower-level hiring through fellowship programs designed for recent graduates and campus recruiting partnerships. However, our findings hold even for management-level employees (levels 3-4), where these initiatives are much less applicable. The persistence of effects at management levels strengthens confidence that the Bias Audit Law, rather than these confounding events, drives our observed results.

We further explore whether the reduction in male hiring at lower to mid-tier positions is concentrated in offices that obtain bias audit. Using the same specification based on reporter status, Table A2 Panel B shows that coefficients on *Treat×Post×Report* are significantly negative for entry-level and mid-level positions while coefficients on *Treat×Post×NonReport* are insignificant, further supporting that bias auditing is the primary driver of the reduction in male hiring at these seniority levels.

Private and Public firms

Previous studies find that 80% of U.S. companies and nearly all Fortune 500 firms use AI-powered hiring software (Wiessner 2024; Dennison 2024). To examine whether our results are driven solely by these firms which face greater reputational and regulatory scrutiny, we conduct subsample analyses based on firm public status.

We classify firms as public if they have a GVKEY identifier provided by Revelio Labs. Table A3 Panel A presents the results. Columns (1) and (2) reports estimates for private and public firms, respectively.[22] For both, the *Treat×Post* interaction terms are statistically significant, indicating that both public and private firms reduce male hiring following the Bias Audit Law. These results suggest that the regulatory impact extends beyond large, publicly traded firms and affects broader segments of the labor market.[23]

We further explore whether the reduction in male hiring at both private and public companies is concentrated among firms that had their algorithm audited. Similarly, we use the specification based on reporter status. As Table A3 Panel B shows, *Treat×Post×Report* is significantly negative in both the private and public company subsamples while *Treat×Post×NonReport* is insignificant, further supporting the interpretation that bias audit compliance is the primary driver of the reduction in male hiring across both firm types.

Bias Audits and Hiring Patterns of Non-White Racial Subgroups

To better understand the impact of Bias Audit Laws on different racial groups, we conduct an additional analysis focusing on non-white subgroups. In the Revelio Labs dataset, race categories are defined as mutually exclusive. We re-estimate our main specification by replacing the dependent variable with the proportion of hires for each subgroup: *BlackProp* (Black hires), *HispanicProp* (Hispanic hires), and *AsianProp* (Asian hires). Table A4 presents the results, with Columns (1)–(3) corresponding to each subgroup. The coefficient on

[22] In Table 7, the number of observations for private firms is smaller than that for public firms due to data limitations in the merging process. We merge datasets from Revelio Labs, the Individual Database and the Job Posting Database, with the Workforce Dynamics Database separately, which provides total employment information. However, as confirmed by a Revelio Labs representative, not all companies in the Individual or Job Posting Databases can be matched to the Workforce Dynamics Database. This issue is especially common among smaller firms, which may have detailed job posting or individual-level data but lack complete employment records in Workforce Dynamics. As a result, when total employment is included as a control, the sample size for private firms (firms without GVKEY) is substantially reduced. If we exclude total employment as a control and do not merge with the Workforce Dynamics Database, the number of observations in the private firm sample becomes roughly four times larger than that in the public firm sample (firms with GVKEY). Importantly, our findings remain robust, suggesting that the main results are not driven by the inclusion of employment controls or the smaller sample.

[23] In untabulated results, we also examine the heterogeneity of the Bias Audit Law's effects by firm size. We find that the reduction in male hiring following the law's enforcement is concentrated among large and medium-sized firms, which are more likely to adopt AI-based hiring practices.

Treat×*Post* is significantly positive in Column (1), indicating an increase in Black hiring, while it is significantly negative in Column (2), suggesting a decline in Hispanic hiring. These findings imply that while the overall share of non-white hires may remain stable, Bias Audit Laws may lead to a redistribution of hiring opportunities among non-white racial groups.

Table A1: Pre-Law Gender Bias in AI Hiring: Evidence from Audit Reporter Offices

This table tests whether AI hiring systems are on average more biased toward males prior to regulatory intervention, using office-quarter observations for the NYC office subsample. We use audit reporter offices as a proxy for confirmed AI hiring adopters and examine whether these offices exhibit higher male hiring rates before the enforcement of the Bias Audit Law. The dependent variable is the proportion of male hires (*MaleProp*). *Report* equals 1 if the office belongs to a firm that has issued a bias audit report, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

	(1)	(2)
DV =	*MaleProp*	*MaleProp*
Report	**0.015***	**0.016***
	(1.744)	**(1.743)**
LogEmpNum	-0.007***	-0.007***
	(-2.980)	(-2.979)
Industry FE	Yes	No
Year-Quarter FE	Yes	No
Industry-Year-Quarter FE	No	Yes
R^2	0.188	0.228
N. of Obs.	7,959	7,333

Table A2: Subsample Analyses: Job Seniority

This table examines how the Bias Audit Law affected male hiring shares by job seniority, using office-quarter observations. The dependent variable is the proportion of male hires (*MaleProp*). *Treat* equals 1 for offices located in New York City, and *Post* equals 1 for the quarter the law takes effect and thereafter. Panel A reports results across three Revelio Labs seniority tiers: Junior (Seniority = 1 or 2), Mid (Seniority = 3, 4, or 5), and Senior (Seniority = 6 or 7), where 1 = entry, 2 = junior, 3 = associate, 4 = manager, 5 = director, 6 = executive, and 7 = senior executive. Panel B further compares the treatment effects between audit reporters and non-reporters within each seniority tier. *Report* equals one if a firm has issued bias audit reports, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. *NonReport* equals one for a firm not classified as a *Report* firm. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Job Seniority and Hiring Diversity

	(1) Junior (Seniority = 1 or 2)	(2) Mid (Seniority = 3 or 4 or 5)	(3) Senior (Seniority = 6 or 7)
DV =		*MaleProp*	
Treat* × *Post	**-0.016****	**-0.020*****	-0.009
	(-2.382)	**(-2.792)**	(-0.591)
LogEmpNum	-0.005**	-0.007**	-0.023**
	(-2.286)	(-2.485)	(-2.381)
Office FE	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes
R^2	0.386	0.371	0.363
N. of Obs.	354,540	254,258	40,228

Panel B: Audit Reporters versus Non-Reporters across Seniority Levels

	(1) Junior (Seniority = 1 or 2)	(2) Mid (Seniority = 3, 4, or 5)	(3) Senior (Seniority = 6 or 7)
DV =		*MaleProp*	
***Treat* × *Post* × *Report* (i)**	**-0.025*****	**-0.030*****	0.002
	(-2.594)	**(-2.992)**	(0.107)
***Treat* × *Post* × *NonReport* (ii)**	-0.008	-0.010	-0.019
	(-0.955)	(-1.038)	(-0.976)
LogEmpNum	-0.005**	-0.007**	-0.023**
	(-2.281)	(-2.478)	(-2.383)
F-statistics for (i) = (ii)	*1.73*	*1.99*	*0.61*
Office FE	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes
R^2	0.386	0.371	0.363
N. of Obs.	354,540	254,258	40,228

Table A3. Subsample Analyses: Private and Public Firms

This table examines the effect of the Bias Audit Law on the proportion of male hires for the private and public firm subsamples, using office-quarter observations. The dependent variable is the proportion of male hires (*MaleProp*). *Treat* equals 1 for offices located in New York City, and *Post* equals 1 for the quarter the law takes effect and thereafter. Panel A reports results on private and public firm subsamples. Column (1) reports estimates for private firms (firms without GVKEY), and Columns (2) for public firms (firms with GVKEY). Panel B further compares treatment effects between audit reporters and non-reporters within each subsample. *Report* equals 1 if a firm has issued a bias audit report, either directly by publishing an audit report or indirectly by adopting any one of the 11 verified algorithms whose vendors have posted a bias audit report. *NonReport* equals one for a firm not classified as a *Report* firm. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

Panel A: Private versus Public Firms

	(1) Private Firms	(2) Public Firms
DV =	*MaleProp*	
Treat* × *Post	**-0.018****	**-0.018*****
	(-1.975)	**(-2.766)**
LogEmpNum	-0.003	-0.006***
	(-1.170)	(-3.376)
Office FE	Yes	Yes
Year-Quarter FE	Yes	Yes
R^2	0.378	0.368
N. of Obs.	136,070	341,505

Panel B: Private versus Public Firms by Reporter Status

	(1) Private Firms	(2) Public Firms
DV =	*MaleProp*	
***Treat* × *Post* × *Report* (i)**	**-0.028****	**-0.028*****
	(-2.344)	**(-2.907)**
***Treat* × *Post* × *NonReport* (ii)**	-0.005	-0.012
	(-0.383)	(-1.439)
LogEmpNum	-0.003	-0.006***
	(-1.167)	(-3.374)
F-statistics for (i) = (ii)	*1.56*	*1.62*
Office FE	Yes	Yes
Year-Quarter FE	Yes	Yes
R^2	0.378	0.368
N. of Obs.	136,070	341,505

Table A4: Heterogeneous Effects on Non-White Subgroup Hiring Shares

This table examines how the Bias Audit Law affects hiring shares for different non-white racial groups, using office-quarter observations. Columns (1)–(3) report results for *BlackProp* (proportion of Black hires), *HispanicProp* (proportion of Hispanic hires), and *AsianProp* (proportion of Asian hires), respectively. The primary explanatory variable is the interaction term between *Treat* and *Post*, where *Treat* equals one for offices located in New York City and *Post* equals one for the quarter the law takes effect and thereafter. Standard errors are clustered at the office level. T-statistics are reported below the coefficients. ***, **, and * correspond to 1%, 5% and 10% significance level. Detailed variable definitions are provided in Appendix Table A.

DV =	(1) *BlackProp*	(2) *HispanicProp*	(3) *AsianProp*
Treat* × *Post	0.008***	-0.007**	0.001
	(2.637)	(-2.348)	(0.169)
LogEmpNum	0.000	0.002**	0.000
	(0.079)	(2.554)	(0.457)
Office FE	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes
R^2	0.316	0.371	0.402
N. of Obs.	477,575	477,575	477,575