FAIRLY

RSP AI Partners - Resume Analytics Platform

New York City Local Law 144 Bias Audit Report

Version 1.0.0 | Last Updated: May 19, 2025

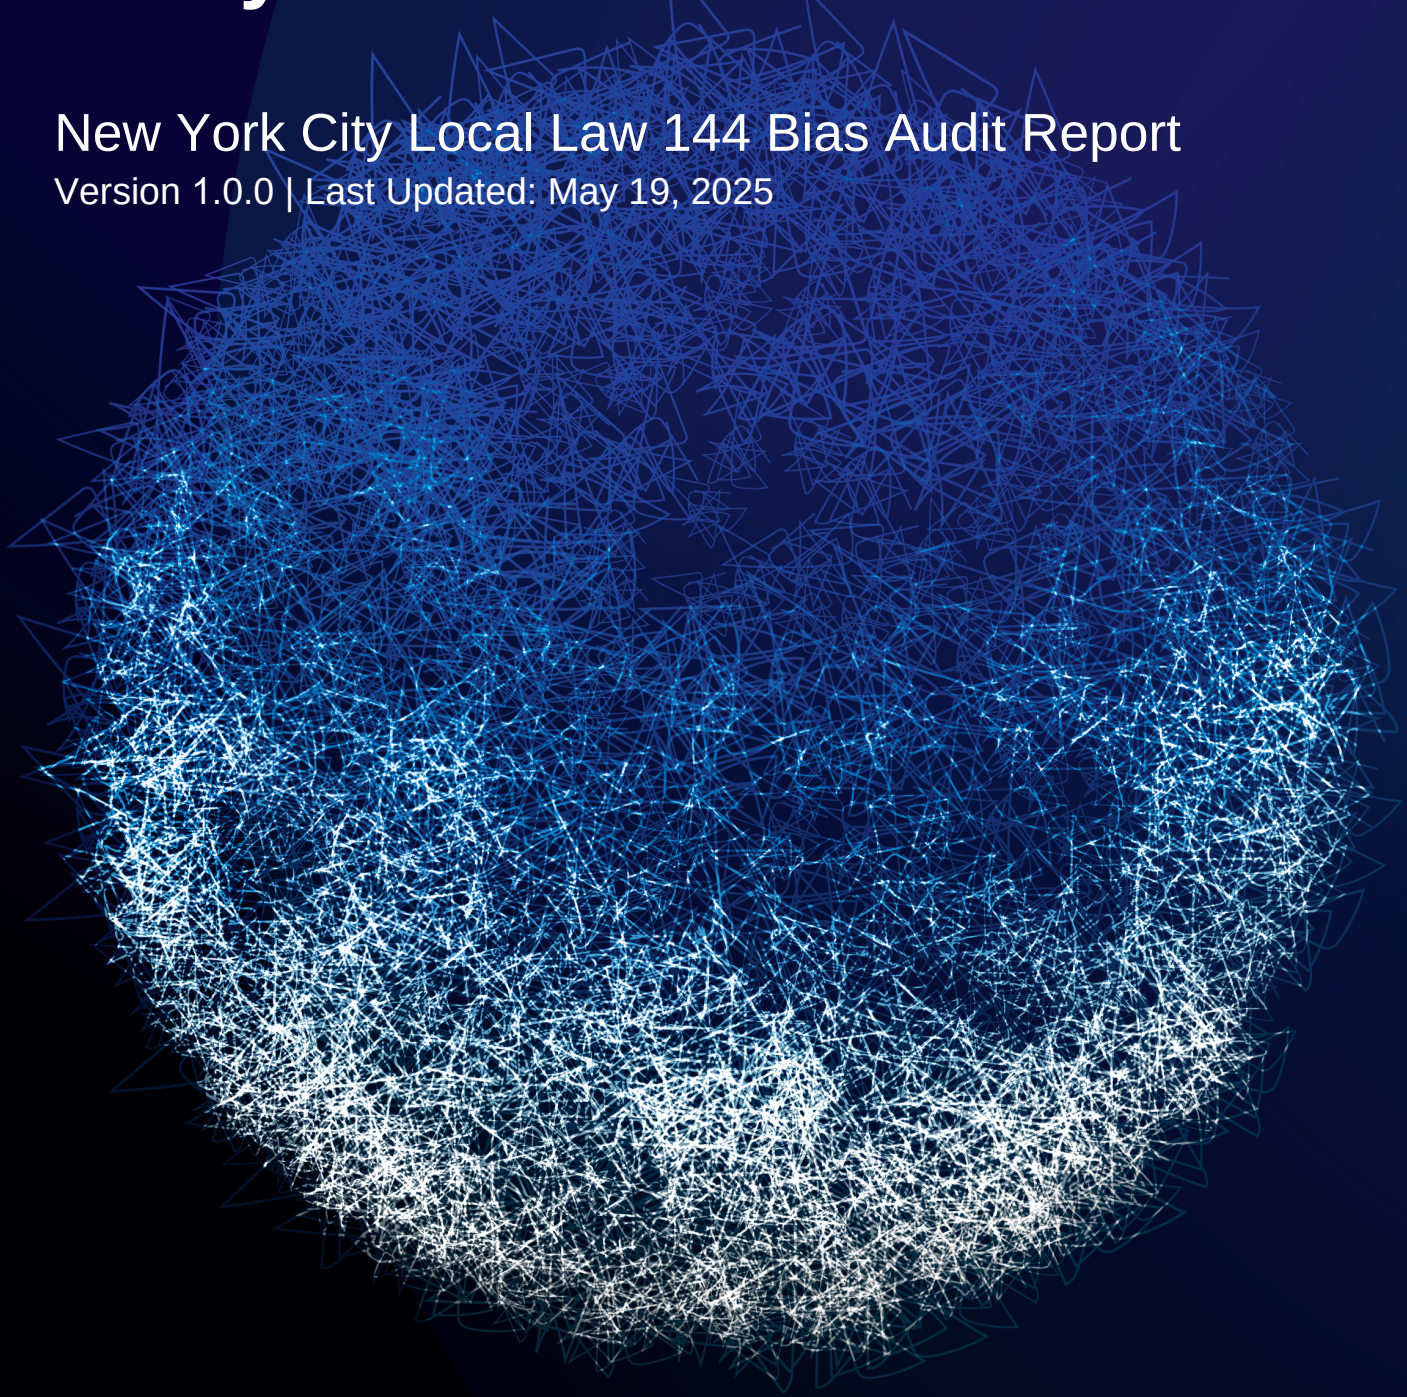

FAIRLY

Compliance Statement

This report has been prepared by Fairly AI, Inc., an independent auditor, in accordance with the requirements of New York City Local Law 144 for Automated Employment Decision Tools (AEDTs). It provides an independent bias audit but does not categorically make a determination as to whether RSP AI Partners's use of Resume Analytics Platform qualifies as an AEDT under the law, nor does it evaluate compliance with the notice or public disclosure requirements of NYC Local Law 144.

Key Findings

Tool Audited: RSP AI Partners - Resume Analytics Platform
Audit Period: March - April 2025
Sample size: 21600
Our analysis found no evidence of disparate impact, as all groups had a selection rate of at least 80% of the most favored group's selection rate across (i) sex, (ii) race/ethnicity, and (iii) all combinations of sex and race/ethnicity, provided the sample size met the minimum threshold for reliable analysis.

Summary

This report summarizes the bias audit results and methodology for Resume Analytics Platform, developed by RSP AI Partners, in compliance with New York City Local Law 144 Bias Audit.

Our audit examines scoring patterns across demographic groups to ensure fair treatment of all applicants. The analysis compares the percentage of candidates scoring above the minimum threshold (30th percentile) across demographic categories. To ensure reliable conclusions, we require sufficient sample sizes that allow us to detect meaningful differences between groups.

The data used in this audit was provided by RSP AI Partners. Fairly AI has conducted the analysis based on this data and has not independently verified its accuracy, completeness,

or representativeness. The findings of this audit are therefore reliant on the integrity of the data as provided by RSP AI Partners.

Tool Description

Description of the tool, including how candidates are scored.

RSP is an AI-driven platform for the analysis of resumes to facilitate comparison of resumes, CVs and other similar biographical information to job, sector or industry criteria as one component in the identification or prioritization of suitable candidates and mutually beneficial matches for employers and individuals interested in pursuing employment opportunities. There are broadly two editions of RSP with different capabilities: The Personal Edition processes is aimed at individuals who wish to analyze their resume either against 50+ industry/sector profiles or against custom-created profiles reflecting a position, job specification or other role description. The Enterprise Edition handles volume processing of resumes in batches of up to 100 resumes. The system allows users to select predefined job role profiles or configure custom profiles based on specific requirements. Input is facilitated through the upload of resumes in PDF format. An advanced AI engine processes each resume across 10 key dimensions relevant to the selected profile or custom criteria. The output includes a comparative analysis of candidates, provided as PDF for single resumes or in CSV and/or PDF format for batches, containing individual dimension and total point scores, and individual PDF reports detailing strengths, areas for improvement, and recommendations for each candidate. These reports support data-informed hiring decisions and team collaboration.

AEDT Classification

This tool may qualify as an Automated Employment Decision Tool under NYC Local Law 144 because it: 1. Uses computational processes derived from artificial intelligence. 2. Substantially assists employment decision-making by providing standardized resume scoring algorithm. 3. Materially impacts natural persons by influencing which candidates advance in the hiring process

Implementation Process

The implementation process will be specific to clients' specific needs.

Data Collection

How demographic data was collected

Synthetic candidate data was created to reflect a broad and diverse population. The dataset spans multiple job functions and covers a range of seniority levels from entry to senior. Each demographic cohort includes balanced representation across gender and ethnicity, with each intersection represented by 100 candidates. The data also reflects varied educational backgrounds, with candidates drawn from a wide spectrum of institution types, ranging from community colleges to elite universities. This structured diversity ensures high-confidence testing for fairness across key demographic dimensions.

Audit Results

Candidate Scoring Methodology Analysis

Each applicant receives an associated score ranging from 0 to 1000 (comprised of 10 individual dimension scores ranging from 0-100) based on the analysis of the candidate’s resume. Since all candidates are shown to the recruiter and there is not a designated pass/fail cutoff, the scoring rate methodology is in alignment with NYC Local Law 144's recommendation

Impact Ratio Analysis: Sex

Sex	# of Applicants	Scoring Rate	Impact Ratio
Male	7,200	67.3%	0.98
Female	7,200	68.4%	1.00
Prefer not to answer	7,200	67.6%	0.99

Impact Ratio Analysis: Race/Ethnicity

Race/Ethnicity	# of Applicants	Scoring Rate	Impact Ratio
Asian (Not Hispanic or Latino)	2,700	67.5%	0.98
White (Not Hispanic or Latino)	2,700	67.5%	0.98
Black or African American (Not Hispanic or Latino)	2,700	67.4%	0.97
Hispanic or Latino	2,700	69.2%	1.00
Two or more races (Not Hispanic or Latino)	2,700	67.9%	0.98
American Indian or Alaska Native (Not Hispanic or Latino)	2,700	67.2%	0.97
Native Hawaiian or Other Pacific Islander (Not Hispanic or Latino)	2,700	67.4%	0.97
Prefer not to answer	2,700	68.0%	0.98

Impact Ratio Analysis: Intersectional

Race/Ethnicity	Sex	# of Applicants	Scoring Rate	Impact Ratio
Asian (Not Hispanic or Latino)	Male	900	67.6%	0.95
Asian (Not Hispanic or Latino)	Female	900	67.8%	0.95
Asian (Not Hispanic or Latino)	Prefer not to answer	900	67.1%	0.94
White (Not Hispanic or Latino)	Male	900	66.9%	0.94
White (Not Hispanic or Latino)	Female	900	68.6%	0.96
White (Not Hispanic or Latino)	Prefer not to answer	900	67.1%	0.94
Black or African American (Not Hispanic or Latino)	Male	900	67.6%	0.95
Black or African American (Not Hispanic or Latino)	Female	900	67.7%	0.95
Black or African American (Not Hispanic or Latino)	Prefer not to answer	900	67.0%	0.94
Hispanic or Latino	Male	900	67.8%	0.95
Hispanic or Latino	Female	900	71.7%	1.00
Hispanic or Latino	Prefer not to answer	900	68.7%	0.97
Two or more races (Not Hispanic or Latino)	Male	900	67.3%	0.95

Race/Ethnicity	Sex	# of Applicants	Scoring Rate	Impact Ratio
Two or more races (Not Hispanic or Latino)	Female	900	68.2%	0.96
Two or more races (Not Hispanic or Latino)	Prefer not to answer	900	68.2%	0.96
American Indian or Alaska Native (Not Hispanic or Latino)	Male	900	67.0%	0.94
American Indian or Alaska Native (Not Hispanic or Latino)	Female	900	67.6%	0.95
American Indian or Alaska Native (Not Hispanic or Latino)	Prefer not to answer	900	67.1%	0.94
Native Hawaiian or Other Pacific Islander (Not Hispanic or Latino)	Male	900	67.3%	0.95
Native Hawaiian or Other Pacific Islander (Not Hispanic or Latino)	Female	900	67.4%	0.95
Native Hawaiian or Other Pacific Islander (Not Hispanic or Latino)	Prefer not to answer	900	67.3%	0.95
Prefer not to answer	Male	900	67.0%	0.94
Prefer not to answer	Female	900	68.9%	0.97
Prefer not to answer	Prefer not to answer	900	68.0%	0.96

Appendix: Statistical Analysis and Methodology

This section explains the three essential testing concepts used in the audit: Power Analysis, Impact Ratio, and Statistical significance. These measures help determine: Whether the AEDT treats demographic groups fairly and whether the results of our audit are statistically sound and reliable.

Power Analysis

Before testing, we conducted a power analysis to confirm that our methods could reliably detect disparities if they exist.

1. We aimed to detect a 20% difference in outcomes between groups.
2. The analysis was calibrated for 80% power, meaning we had an 8 in 10 chance of detecting a real issue.
3. We controlled the risk of false positives at 5% (standard industry practice), meaning there's only a 1 in 20 chance that a difference flagged as meaningful is actually due to random variation.
4. Given an overall outcome rate of 67.8%, this resulted in a minimum requirement of 96 candidates per group.
5. This threshold represents approximately 0.4% of the full dataset (n = 21600).

Impact Ratio

We used impact ratios to measure fairness across demographic groups.

1. The impact ratio compares each group's pass or scoring rate to that of the group with the highest rate.
2. A ratio below 0.80 may suggest a potential fairness concern (based on the "four-fifths rule").
3. A ratio below 0.60 is considered a high-risk indicator and should be reviewed immediately.

Statistical Significance

We use statistical tests to determine if observed differences between groups are meaningful or could be due to random chance.

1. Statistical significance tells us if an observed difference is larger than what we would expect from random variation.
2. For each comparison, we calculate a p-value, which is the probability of seeing a difference as large or larger than the one we observed if there were actually no difference.
3. A p-value below 0.05 (5%) suggests the difference is statistically significant and less likely to be due to random chance.

By combining these three measures, we ensure our audit provides a comprehensive and reliable

assessment of potential disparate impact.

About New York City Local Law 144 Bias Audits

New York City Local Law 144 mandates that employers and employment agencies conduct independent bias audits on automated employment decision tools (AEDTs) before using them in hiring or promotion decisions. The audits must assess whether the AEDT exhibits disparate impact based on sex, race/ethnicity, or other protected characteristics by analyzing selection rates. Employers must also provide public disclosures about the audit results and notify candidates when an AEDT is used in their evaluation. Compliance with Local Law 144 ensures greater transparency and fairness in hiring practices.

About Fairly AI

Fairly AI is a leader in AI Trust, Risk, and Security Management, specializing in independent qualitative assessments and quantitative testing for AI models and LLM applications. With proven expertise in financial services, Fairly AI extends established Model Risk Management governance processes—including stress testing, scenario analysis, and benchmarking—to AI, ensuring safe, secure, and compliant adoption. Our work with institutions like the Vector Institute and Partnership on AI demonstrates our commitment to AI ethics and governance. As NYC Local Law 144 requires rigorous bias audits of automated employment decision tools, Fairly AI's deep technical expertise, regulatory knowledge, and independent testing capabilities make us a highly qualified auditor. Fairly AI is SOC2 Type 2 and ISO/IEC 42001 certified.