

Bias Audit for Ashby's Criteria Evaluation Model

Ashby



Table of contents

Table of contents	2
Bias evaluation for Ashby’s Criteria Evaluation model	2
Introduction and executive summary	2
Background	2
About NYC LL144	2
Bias audit requirements	3
System description	3
Bias evaluation results	4
Data used	4
Univariate Categories	5
Results by gender	5
Results by race	5
Intersectional Categories	6
Results by race and gender	6
About FairNow	6

Bias evaluation for Ashby’s Criteria Evaluation model

Report prepared by FairNow on August 23, 2024

Introduction and executive summary

This report presents a disparate impact analysis of Ashby’s Criteria Evaluation model in alignment with the specifications for New York City’s Local Law 144 (referred to as NYC LL144, or the NYC AEDT law). Ashby’s Criteria Evaluation model evaluates job applicant resumes against user-defined job criteria for a given role, and provides a determination of whether the applicant meets each criterion. Users review Ashby model outputs for a given applicant at the criterion level.

The results of this evaluation did not indicate evidence of disparate impact, as there were no groups for which the selection rate was less than 80% of the selection rate of the most favored group when analyzing race, gender or combinations thereof.

Background

About NYC LL144

NYC LL144 entered effect on July 5th, 2023. The law prohibits employers and employment agencies from using Automated Employment Decision Tools (AEDTs) to make or substantially assist in hiring or promotion decisions without meeting specified requirements, detailed below:

1. **Annual bias audit:** Employers must engage an independent auditor to conduct a bias audit of any in-scope AEDTs to evaluate candidate pass rates along the lines of gender, race and intersectional pairs of the two.
2. **Notice to candidates:** Employers must provide candidates with at least 10 business days notice before administering the AEDT, and include instructions on how to request an alternative assessment method if one is available.
3. **Published results:** Employers must publish the date and a summary of the results of the most recent bias audit on the employment section of their website.

This report represents an independent bias audit in alignment with the requirements of NYC LL144. It does not entail an evaluation of whether usage of Ashby's Criteria Evaluation model constitutes an AEDT under NYC LL144 specifications, and it does not include a review of the notice or published result specifications under NYC LL144.

Bias audit requirements

A bias audit as defined by NYC LL144 requires the following:

- The assessment includes a review of gender, race, and intersectional pairs of the two.
- A 'selection' or 'scoring' rate is calculated for each category.
- An impact ratio is calculated for each category.
- Counts are included for applicants with unknown gender or race.

The evaluation outlined in this document has been conducted in accordance with the NYC LL144 bias audit requirements.

System description

The intent of Ashby's Criteria Evaluation model is to aid recruiting and hiring teams in assessing candidate qualifications against user-defined job criteria. Users are responsible for uploading job postings, and specifying the set of job criteria that are relevant for each posting. Ashby leverages an algorithm to parse applicant resumes and provide a determination of whether each applicant meets each of the specified job criteria. Applicant resumes are evaluated so that they will receive a designation of "meets", "does not meet" or "uncertain" against each job criterion. The algorithm generates an "uncertain" result when Ashby's systems are unable to make a confident determination of whether a candidate does or does not meet the criterion based on the existing resume text.

Ashby's product does not aggregate individual applicant-criteria outcomes into applicant-level scores for rank-ordering or evaluation purposes. In alignment with product usage, this audit evaluates

scoring at the applicant-criteria level, where “meets” represents a passing score. The population of applicant-criteria flagged by Ashby’s model as “uncertain” (11.4% of applicant-criteria) were excluded from analysis for the purposes of this audit.

Bias evaluation results

Data used

This evaluation includes historical data from applicants assessed by Ashby’s Criteria Evaluation model during the period of January 1st, 2024 to August 12th, 2024, covering 123,610 applicant-criteria for Ashby job postings across the US and Europe. This data represents all applicant-criteria evaluated by the model during this period.

The data sample leveraged for this audit encompasses a mix of historical data from the actual usage of the Criteria Evaluation model for Ashby postings in addition to applicant-criteria that were retroactively scored after the model was deployed. The Ashby Criteria Evaluation model was launched for early access on July 1st 2024 (with general access anticipated for September 10th, 2024); retroactive scores on actual applicant-criteria from earlier in 2024 were included to supplement the dataset to ensure sample robustness. To retroactively score the model, the Ashby recruiting team reviewed historical job postings from the sample period, defined each posting’s job criteria, and ran the model against historical applications for those job postings to generate applicant-criteria values.

Demographic data was provided from a survey produced by Ashby where candidates were asked to self-select their race/ethnicity and gender. Demographic data is missing on candidates who chose a “decline to specify” option in the survey or did not respond to the question. Candidate demographic information was not imputed or inferred in any way.

Gender was represented as one of the following values:

- Female
- Male
- Unknown (not included in the analysis)
- Decline to specify (not included in the analysis)

Race/Ethnicity was represented as one of the following values:

- Asian
- Black
- Hispanic/Latino
- Two or More Races
- White
- Unknown (not included in the analysis)
- Decline to specify (not included in the analysis)

Univariate Categories

Results by gender

Gender	# of Applicant-criteria	# Selected	Selection Rate	Impact Ratio
Female	45,855	30,120	65.7%	94.1%
Male	72,124	50,358	69.8%	100.0%

There were 5,631 applicant-criteria for whom gender was not known. This data was not included in the above table.

Results by race

Race	# of Applicant-criteria	# Selected	Selection Rate	Impact Ratio
Asian	34,259	22,057	64.4%	90.0%
Black	16,653	10,912	65.5%	91.6%
Hispanic/Latino	14,025	9,388	66.9%	93.6%
Two or More Races	25	15	60.0%	N/A
White	39,114	27,983	71.5%	100.0%

"N/A" refers to categories that comprised less than 2% of the total. These groups are excluded from the disparate impact analysis.

There were 19,534 applicant-criteria for which the applicant's race/ethnicity was not known. This data was not included in the above table.

Intersectional Categories

Results by race and gender

Race	Gender	# of Applicant-criteria	# Selected	Selection Rate	Impact Ratio
Asian	Female	11,861	7,537	63.5%	86.4%
Asian	Male	20,108	13,014	64.7%	88.0%
Black	Female	7,935	5,085	64.1%	87.1%
Black	Male	7,627	5,188	68.0%	92.5%
Hispanic/Latino	Female	5,544	3,575	64.5%	87.7%
Hispanic/Latino	Male	8,285	5,693	68.7%	93.4%
Two or More Races	Male	25	15	60.0%	N/A
White	Female	13,252	9,014	68.0%	92.5%
White	Male	24,728	18,190	73.6%	100.0%

“N/A” refers to categories that comprised less than 2% of the total. These groups are excluded from the disparate impact analysis.

There were 24,245 applicant-criteria for which the applicant’s gender or race/ethnicity were not known. This data was not included in the above table.

About FairNow

FairNow is an organization dedicated to helping companies leverage AI in a responsible, fair, and well-managed way. FairNow is an independent auditor in alignment with the specifications of New York City Local Law 144.