

Bias, Fairness & Validation Report

Synthetic adverse-impact and validation assessment for EMP-AI-01 ResumeRank before production deployment.

SIMULATED OPERATING DOCUMENT

Prepared for a fictional employer using synthetic systems and synthetic applicant data. This document is a governance demonstration, not legal advice.

SHOWCASE ARTIFACT

Synthetic data are used solely to demonstrate the analysis. The report deliberately separates outcome disparity, statistical uncertainty, job-relatedness evidence, and governance decision.

1. Test scope and methodology

Population: 1,000 synthetic applicants to the Customer Operations Specialist job family. The simulated rule advances approximately the top 36% of model-ranked applicants to recruiter review. Demographic fields are assumed to exist in a segregated audit dataset and are not model inputs.

- Calculate selection rate = selected / applicants for each group.
- Calculate impact ratio = group selection rate / highest observed comparison-group selection rate.
- Flag ratios below 0.80 as a preliminary adverse-impact signal, not a legal conclusion.
- Review statistical uncertainty using chi-square / pairwise tests where sample sizes support them; do not treat a p-value as proof of discrimination or fairness.
- Investigate feature mechanisms and test a less-discriminatory alternative rather than merely documenting the baseline.
- Evaluate job-relatedness and business-necessity evidence separately from disparity metrics.

Authority notes: R2, R7.

2. Baseline race / ethnicity results

Group	Applicants	Advanced	Selection rate	Impact ratio	Screen
White	400	160	40.0%	1.00	Reference / highest
Black	250	75	30.0%	0.75	FLAG: below 0.80
Hispanic / Latino	200	68	34.0%	0.85	Monitor
Asian	130	52	40.0%	1.00	Reference / highest
Other / Multiracial	20	7	35.0%	0.88	Small N; interpret cautiously

Synthetic omnibus chi-square: $\chi^2 = 7.91$, $df = 4$, $p = 0.095$. Pairwise Black vs White Fisher exact test: $p \approx 0.012$. These values illustrate why the 0.80 screen and statistical testing should be considered together rather than substituting one for the other.

3. Baseline sex and intersectional spot check

Group	Applicants	Advanced	Selection rate	Impact ratio vs highest
Men	520	209	40.2%	1.00
Women	460	149	32.4%	0.81
Nonbinary / another reported category	20	4	20.0%	0.50 (very small N)

Intersectional spot check: Black women 32/125 = 25.6%; White men 95/220 = 43.2%; ratio ≈ 0.59. This is a governance trigger for deeper analysis, but the synthetic subgroup size and multiple-comparison context limit inference.

4. Root-cause investigation

Finding	Evidence / reasoning	Governance conclusion
Career-continuity feature contributes to score gap	Feature ablation shows gap narrows when employment-gap length and job-change count are removed	Feature is not sufficiently tied to essential job requirements; remove.
Current-title prestige contributes to score gap	Candidates with equivalent skill evidence receive different scores based on title normalization	Replace prestige signal with explicit skill/level evidence.
Résumé parser failures are uneven by format	Multi-column and accessibility-generated formats show higher extraction error in QA sample	Add plain-text alternative, parser quality check, and correction route.
Vendor "bias-free" claim is insufficient	Claim does not establish disability accessibility, job relevance, or client-population outcomes	Require evidence; conduct employer-side assessment.

Authority notes: R3.

5. Less-discriminatory alternative test

Mitigated configuration removes employment-gap length, job-change count, graduation year, address/ZIP, name/entity prestige, and any inferred personality/affect signals. Current title is normalized to a role/skill taxonomy with capped weight.

Group	Applicants	Advanced	Selection rate	Impact ratio
White	400	153	38.3%	0.99
Black	250	87	34.8%	0.90
Hispanic / Latino	200	72	36.0%	0.94
Asian	130	50	38.5%	1.00
Other / Multiracial	20	7	35.0%	0.91

Synthetic mitigated omnibus chi-square: $\chi^2 \approx 1.02$, $df = 4$, $p \approx 0.906$. The practical group-rate spread also narrows substantially. This does not establish legal compliance or validity; it shows that a lower-disparity configuration is available and therefore must be considered in the governance decision.

6. Job-relatedness and validation status

Evidence item	Status	Assessment
Current job analysis	Partial	Competency profile exists but requires HR validation for this exact job family.
Construct / content linkage	Partial	Skill and experience criteria are plausibly related; continuity/title prestige are not adequately supported.
Criterion-related validation	Not complete	No client-specific study demonstrates that the ranking score predicts a defensible job-performance criterion.
Accessibility validity	Not complete	No evidence yet that parser/assessment accuracy is comparable for candidates using alternative formats or accommodations.
Transportability of vendor validation	Insufficient	Vendor aggregate study is not enough to establish client-specific use validity.

7. Decision

NO-GO FOR BASELINE

Do not launch the baseline model. The mitigated model may proceed to a limited pilot only after job-analysis validation, accessibility testing, human-review configuration, legal review, notice/recourse implementation, and monitoring are complete.

8. Statistical and legal limitations

- The dataset and results are synthetic and do not describe a real employer or vendor.
- The four-fifths rule is used as a screening heuristic; it is not a complete legal or statistical test.
- Small group sizes can produce unstable rates; larger samples can make small effects statistically significant.
- Statistical disparity does not by itself establish causation, intent, business necessity, or legality.
- Protected-class data collection and use must be reviewed for lawful basis, notice, segregation, access, and retention.
- A bias audit does not replace disability/accessibility review, job-validation evidence, meaningful human oversight, or candidate recourse.

9. Independent-review status and evidence package

Item	Status
Independent external bias audit	Not performed - this is a synthetic portfolio assessment.
Internal validation review	Demonstrated through documented methodology, feature analysis, and simulated re-test.
Evidence retained	Test population definition; group-rate tables; statistical notes; feature decisions; alternative test; AIA decision; sign-off conditions.
Production reliance	None. A real employer would require client-specific legal, validation, accessibility, privacy, and operational evidence before use.