Available on: https://jmrr.ep-journals.org/ E-ISSN: 3051-3588, P-ISSN: 3051-357X

Journal of Management Research and Review

Volume: 02 Issue: 08 August-2026

Artificial Intelligence in Recruitment: A Review of Efficiency Claims, Bias Risks, And the Governance Imperative for HR Practice

Mehmoona Akram
MSc International Business | Ulster University Birmingham | CIPD Level 5 (In Progress) ORCID I'd: 0009-0006-5001-4381

Syeda Fatima Hussain
Ulster University Birmingham

Raja Muneeb Anwar
Ulster University Birmingham

Published Date: August 19, 2026 **Article DOI: 10.65150/EP-jmrr/V2E8/2026-10**

ABSTRACT: Artificial intelligence has diffused rapidly through the recruitment function, from automated résumé screening and chatbot-led candidate engagement to algorithmic assessment and predictive analytics, and vendors and adopters advance strong claims about efficiency and quality-of-hire gains. Simultaneously, high-profile failures have made algorithmic hiring a focal case in debates about automated discrimination and its regulation. This paper reviews the multidisciplinary literature on AI in recruitment — spanning human resource management, information systems, computer science research on algorithmic fairness, and the emerging regulatory scholarship — to assess what is credibly known about its benefits, its risks, and the conditions that separate the two. The review finds robust evidence for process-efficiency gains but thin and mixed evidence for quality-of-hire improvement; a well-established taxonomy of bias mechanisms (training data bias, proxy discrimination, and feedback loops) with documented instances in deployed systems; and an emerging governance literature converging on auditability, human oversight, and outcome monitoring as the practices that condition whether adoption helps or harms diversity outcomes. The paper develops a governance-centred framework for HR practice, maps it against incoming regulation including the EU AI Act's classification of employment AI as high-risk, and sets out a research agenda focused on the gap between vendor claims and independently verifiable outcomes.

KEYWORDS: artificial intelligence, recruitment, algorithmic bias, HR analytics, talent acquisition, AI governance, EU AI Act, literature review

1. INTRODUCTION

Recruitment has become one of the most intensively automated functions in management. Surveys of HR practitioners report that a large and growing majority of medium and large organisations now use some form of AI or algorithmic automation in hiring — screening résumés, ranking candidates, scheduling and conducting structured video assessments, and predicting performance or attrition from application data (SHRM, 2023). The commercial logic is straightforward: recruitment is high-volume, text-heavy, and repetitive, precisely the profile on which machine learning tools promise the largest efficiency returns.

The same features make hiring a uniquely sensitive domain for automation. Employment decisions allocate life chances; they are legally protected against discrimination in most jurisdictions; and the historical data on which hiring algorithms learn and encode the discriminatory patterns of past human decisions. The best-known cautionary case — a major technology firm's abandoned experimental screening tool that learned to penalise indicators of female applicants — has become shorthand for a now well-theorised class of failures (Dastin, 2018; Barocas & Selbst, 2016). Regulators have responded: New York City's Local Law 144 mandates bias audits of automated employment decision tools, and the EU AI Act classifies employment-related AI as high-risk, attaching conformity, transparency, and oversight obligations.

HR practice thus faces a genuine governance problem: the tools are diffusing faster than the evidence base or the professional norms for supervising them. This review takes stock. Section 2 maps the technology landscape. Section 3 assesses the evidence on efficiency and quality claims. Section 4 reviews the bias literature. Section 5 synthesises the emerging governance scholarship and regulation into a framework for practice. Section 6 concludes with a research agenda.

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

2. THE TECHNOLOGY LANDSCAPE

AI recruitment tools operate at every funnel stage. Sourcing tools search professional networks and databases to identify passive candidates. Screening tools parse résumés and applications, matching them to job requirements with natural language processing and ranking candidates by predicted fit. Engagement tools — conversational chatbots — handle candidate queries, pre-screening questions, and scheduling. Assessment tools evaluate structured interviews, work samples, and gamified psychometrics, in some products analysing video and audio features. Prediction tools model performance, tenure, and offer acceptance from historical employee data, feeding both selection and workforce planning. The generative-AI wave since 2022 has added a further layer: large language models drafting job descriptions and outreach, summarising candidate files, and — on the candidate side — mass-producing tailored applications, an arms-race dynamic whose net effect on signal quality in the hiring funnel is an open question (Tambe et al., 2019; Cappelli & Tavis, 2018).

3. THE EVIDENCE ON EFFICIENCY AND QUALITY CLAIMS

3.1 Process Efficiency

The efficiency evidence is the strongest part of the record. Case studies and practitioner surveys consistently document reductions in time-to-hire and cost-per-hire following screening automation, concentrated at the initial review stage where recruiter hours per hire are highest, and reallocation of recruiter time toward relationship and advisory work (Deloitte, 2021). These findings are plausible and repeated, though it should be noted that most originate from vendors and adopting firms rather than independent evaluation, and survivorship — failed implementations are rarely written up — likely inflates the published average.

3.2 Quality of Hire

The quality evidence is thinner and more interesting. The strongest independent study, Cowgill's (2019) analysis of a résumé-screening algorithm deployed in a large firm, found the machine outperformed human screeners in identifying candidates who would pass interviews and accept offers, and — notably — was less biased than the humans it replaced, favouring non-traditional candidates the human process screened out. Against this, several studies and audits find algorithmic assessments adding little predictive validity over structured traditional methods, and the meta-analytic tradition in personnel psychology is a sobering benchmark: decades of research show that even the best selection instruments predict performance only moderately, a ceiling that machine learning on the same criterion data cannot simply wish away (Schmidt & Hunter, 1998). The honest summary is that AI reliably accelerates hiring; whether it improves recruitment depends on the tool, the baseline, and the deployment — and independent evidence remains scarce relative to the scale of adoption.

4. BIAS: MECHANISMS AND EVIDENCE

The algorithmic fairness literature has converged on a three-mechanism taxonomy directly applicable to hiring. Training-data bias: models learn from historical decisions; where past hiring disadvantaged a group, the model reproduces the disadvantage as a pattern, laundering it through apparent objectivity (Barocas & Selbst, 2016). Proxy discrimination: removing protected attributes does not remove their correlates — postcode, education pathway, career-gap patterns, even lexical style — so facially neutral features reconstruct the protected attribute; this is why 'we don't use gender' is not a defence. Feedback loops: biased selection produces unrepresentative employee data, which trains the next model, entrenching the original skew (O'Neil, 2016).

Documented instances span the funnel: the abandoned screening tool noted above; audit studies finding advertisement delivery algorithms skewing job-ad exposure by gender and age; and independent audits of assessment vendors reporting disparate scoring under conditions such as non-native accents or atypical backgrounds. Two points from this literature deserve emphasis for HR audiences. First, fairness has multiple, mathematically incompatible definitions — demographic parity, equalised odds, calibration — so 'the tool is bias-tested' is meaningless without asking against which definition and threshold (Raghavan et al., 2020). Second, the comparison that matters for practice is not algorithm versus perfection but algorithm versus the human process it replaces, which decades of audit-study research show to be substantially biased; a well-governed algorithm can be an improvement on that baseline, and a poorly governed one an industrialisation of it. Governance, not the technology itself, determines which.

5. A GOVERNANCE FRAMEWORK FOR HR PRACTICE

The emerging governance scholarship and the regulatory instruments now in force converge on a recognisable set of practices, which this review organises into a four-layer framework. Procurement diligence: demanding validation evidence, adverse-impact statistics by subgroup, and audit access before adoption — treating vendor claims as claims. Deployment controls: human review of algorithmic rejections at defined risk tiers; candidate transparency about automated processing (a legal requirement under the EU AI Act's high-risk regime and NYC Local Law 144); and documented override authority so that accountability remains locatable in a person. Outcome monitoring: continuous measurement of selection rates and progression by protected group against pre-committed thresholds, because bias in adaptive systems is a moving property, not a one-time certification. Capability building: the framework's binding constraint is professional competence — HR practitioners able to interrogate a validation study, read an adverse-impact ratio, and exercise informed override judgement — which positions AI governance as a core content area for professional bodies such as the CIPD rather than a specialist afterthought (Tursunbayeva et al., 2022).

The framework aligns with, but exceeds, current legal minima: regulation establishes floors (audits, notices, conformity assessment), while the reputational and litigation exposure of hiring discrimination — and the diversity commitments most large employers have made voluntarily — argue for treating the monitoring layer as continuous management practice rather than compliance ritual. It also implies a division of labour: fairness metrics and audits are necessary but technical; deciding which fairness definition an organisation will stand behind is a values judgement that belongs with HR and leadership, not with the vendor's data scientists.

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

6. CONCLUSION AND RESEARCH AGENDA

The literature reviewed here supports a deliberately unexciting conclusion: AI in recruitment is neither the objectivity machine of vendor marketing nor the inevitable discrimination engine of the darkest commentary, but a set of tools whose effects are conditional on governance — and the governance evidence, unlike the efficiency evidence, is still thin. The research priorities follow. First, independent outcome evaluation: the field needs studies of quality-of-hire and diversity outcomes conducted by researchers without commercial stakes, across the many organisations now generating exactly the data required. Second, the regulatory natural experiments: NYC's audit mandate and the EU AI Act's phased application create before-and-after variation that can identify whether mandated auditing changes tool design and hiring outcomes, not merely paperwork. Third, the generative-AI funnel: the collision of LLM-assisted applications with LLM-assisted screening is reshaping the information content of the application itself, and hiring research has barely begun to measure it. Fourth, practitioner capability: what training actually equips HR professionals to govern these systems is an empirical question the profession has so far answered only with syllabi. For a function whose decisions allocate livelihoods, the gap between adoption speed and evidence quality is itself the central finding of this review — and closing it is a task for exactly the kind of practice-grounded scholarship that HR research exists to provide.

REFERENCES

1) Barocas, S., & Selbst, A.D. (2016). Big data's disparate impact. California Law Review, 104(3), 671–732.
2) Cappelli, P., & Tavis, A. (2018). HR goes agile. Harvard Business Review, 96(2), 46–52.
3) Cowgill, B. (2019). Bias and productivity in humans and algorithms: Theory and evidence from résumé screening. Columbia Business School Working Paper.
4) Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters, 10 October 2018.
5) Deloitte. (2021). 2021 Deloitte global human capital trends. Deloitte Insights.
6) O'Neil, C. (2016). Weapons of math destruction: How big data increases inequality and threatens democracy. Crown Publishers.
7) Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 469–481.
8) Schmidt, F.L., & Hunter, J.E. (1998). The validity and utility of selection methods in personnel psychology: Practical and theoretical implications of 85 years of research findings. Psychological Bulletin, 124(2), 262–274.
9) SHRM. (2023). The state of AI in HR 2023. Society for Human Resource Management.
10) Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resources management: Challenges and a path forward. California Management Review, 61(4), 15–42.
11) Tursunbayeva, A., Pagliari, C., Di Lauro, S., & Antonelli, G. (2022). The ethics of people analytics: Risks, opportunities and recommendations. Personnel Review, 51(3), 900–921.
12) Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine bias. ProPublica, 23 May 2016.
13) Kleinberg, J., Mullainathan, S., & Raghavan, M. (2017). Inherent trade-offs in the fair determination of risk scores. Proceedings of Innovations in Theoretical Computer Science (ITCS).
14) Marler, J.H., & Boudreau, J.W. (2017). An evidence-based review of HR analytics. International Journal of Human Resource Management, 28(1), 3–26.
15) Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. Harvard Journal of Law & Technology, 31(2), 841–887.
16) European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act). Official Journal of the European Union.

License: This is an open access article under the CC BY 4.0 license: https://creativecommons.org/licenses/by/4.0/

Corresponding Author: Mehmoona Akram