Augmentation or Attrition? Task Redistribution After a Generative AI HR Deployment

Rohini Kaninde
Clarke University, Worcester, MA, USA

rohinikaninde95@gmail.com

Abstract: Generative artificial intelligence (AI) is widely expected to reshape human resource management (HRM), yet organization-level evidence on how it redistributes work between people and machines is scarce, and debate is split between optimistic augmentation and pessimistic attrition narratives. This paper asks a narrower, answerable question: after an enterprise generative AI platform is deployed in HR, which tasks move to the machine, which remain human, and how do practitioners experience the change? We report an exploratory pre–post study of a single platform deployed across ten organizations over six months, combining anonymized operational data (more than 6,000 requisitions) with thirty-two practitioner interviews. The design is observational, with no control group, so results are reported as associations rather than causal effects. Time-to-hire fell by a median of 23% across sites (95% CI 20-26%; from about 42 to 32 days), and the volume of routine queries reaching human advisers fell by 72% (95% CI 69-75%); every site moved in the same direction (Wilcoxon signed-rank p = 0.002, matched-pairs rank-biserial r = 1.00), though with only ten sites this is indicative. Both varied across sites and accrued gradually over the six months rather than at switch-on. The platform absorbed high-volume, rule-bound tasks such as screening, scheduling, and policy queries, while judgment- and relationship-based tasks remained human. Practitioners reported reallocating time towards that relational work, alongside concerns about transparency, surveillance, and new supervisory demands. We do not observe long-run employment outcomes: no participating organization reduced HR headcount during the window. Six months is shorter than a typical restructuring cycle, and we cannot establish that augmentation, rather than attrition, is the lasting result. We therefore treat the optimistic reading as a conditional leadership choice, not a finding. Because the platform is a commercial product connected to the authors and the data are vendor-context, we disclose this relationship in full and report the study as early, transparency-qualified evidence. The contribution is empirical: a rare multi-site, mixed-methods account of task redistribution after a generative AI HR deployment, offered with explicit attention to its limits and to the governance obligations, including high-risk AI rules, that such deployments now entail.

Keywords: Generative artificial intelligence, Human resource management, Task redistribution, Workforce augmentation, Mixed-methods study, AI governance, Responsible leadership

1. Introduction

The public release of large language model tools such as ChatGPT in late 2022 moved generative AI from a specialist concern to an everyday workplace technology. HRM is unusually exposed because much HR work is high-volume, language-based, and policy-bound activity that these models handle well (Vrontis et al. 2022; Madanchian 2024).

Early discussion has been dominated by cost. IBM reports that an internal HR assistant handled more than 11.5 million employee interactions in 2024 and contributed to a roughly 40% reduction in HR operational costs over four years (IBM 2025), although this is a single company's self-reported figure. Industry and vendor surveys claim further savings, but those estimates are largely unverified, and we do not rely on them. Such headline numbers say little about what actually happens to the work.

Two questions follow. First, are the gains durable or an inflated phase of a familiar cycle, given that most organizations report little measurable financial impact from AI (MIT NANDA 2025)? Second, and our focus, how is work divided between people and machines? International evidence suggests most jobs are exposed to AI for augmentation rather than full automation (ILO 2025), and management scholarship argues that automation and augmentation are interdependent rather than opposed (Raisch and Krakowski 2021); the augmentation idea is itself well established, not new (Davenport and Kirby 2016; Daugherty and Wilson 2018).

Rather than assert which narrative is correct, we ask an answerable, organization-level question and report early evidence. Drawing on the deployment of one enterprise generative AI platform across ten organizations, we combine operational data with practitioner interviews to show which tasks move to the machine and which remain human, and how practitioners experience the shift. Our contribution is primarily empirical: organization-level, multi-site, mixed-method evidence of task redistribution, an area where measured studies are scarce (Tambe, Cappelli, and Yakubovich 2019); we also use that evidence to test and qualify the automation-augmentation paradox at the organizational level (Section 5). Because the platform is a commercial product connected to the authors, we disclose this relationship in full (see Conflict of interest) and

scope our claims accordingly. The paper reviews the literature and theory, sets out the method, reports findings, and discusses implications for HR leadership, AI governance, and the management of change.

2. Literature Review and Theoretical Framing

2.1 Generative AI in Human Resource Management

Research on AI in HRM has grown quickly. A systematic review by Vrontis et al (2022) found that intelligent technologies are reshaping recruitment, training, and performance management. Madanchian (2024) maps tools across the employee lifecycle; Venugopal et al (2024) apply topic modeling to workforce planning; and Rukadikar, Khandelwal, and Warrier (2025) trace a two-decade shift towards AI-supported hiring. The recruitment stage has the most empirical attention: Abdelhay et al (2025), in a survey of 469 professionals, link generative AI to recruitment efficiency, an effect mediated by automation level and moderated by organisational size; Mamadjanova (2025) confirms in a review that screening and chatbot tools reduce time-to-hire. Most of this work, however, is conceptual, survey-based, or single-site; organization-level evidence of how tasks are actually redistributed, and what happens to the staff affected, remains scarce (Tambe, Cappelli, and Yakubovich 2019; Cedro 2025).

2.2 Theoretical Framing

We read the deployment through four complementary lenses, because no single one suffices. At the macro level, the task-based framework of Acemoglu and Restrepo (2019) distinguishes a displacement effect, as machines take over tasks, from a reinstatement effect, as new tasks emerge in which labor holds an advantage; whether labor demand falls depends on which dominates. To explain organization-level change, the AMO framework (Appelbaum et al 2000) and Job Demands–Resources and work-design theory (Bakker and Demerouti 2007) frame how removing routine load can alter the content and quality of HR roles. Sociotechnical systems theory (Trist and Bamforth 1951) anticipates that value comes from jointly redesigning technology and work, not from technology alone. Technology acceptance theory (UTAUT; Venkatesh et al 2003) frames the trust and adoption dynamics we observe. Across these lenses, augmentation and automation are interdependent rather than alternatives (Raisch and Krakowski 2021), so our contribution is empirical rather than a claim to conceptual novelty.

2.3 Critical Perspectives and Research Questions

A critical literature warns against optimistic readings: algorithmic management can intensify work, enable surveillance, and deskill roles even as it removes drudgery (Kellogg, Valentine, and Christin 2020). We therefore treat practitioner accounts sceptically and actively seek disconfirming evidence. Against this backdrop, we ask three exploratory questions:

RQ1: What is the association between deploying the platform and HR efficiency?

RQ2: How is the division of labor between AI and HR staff reconfigured?

RQ3: What are the implications for HR leadership, AI governance, and workforce strategy?

3. Methodology

3.1 Design and Setting

This is an exploratory, pre–post study. It has no control group or counterfactual, so we report associations rather than causal effects and interpret them in light of the threats to validity outlined below. The setting is the rollout of an enterprise generative AI HR platform (anonymized for double-blind review) across 10 organizations spanning technology, financial services, healthcare, the public sector, retail, manufacturing, and other sectors, with employee counts ranging from about 1,000 to more than 20,000 (Table 2). Technically, the platform combines a large language model with retrieval-augmented generation over each organization's HR policy corpus., and a multi-agent layer that routes requests; Recruiters and advisers act as human-in-the-loop reviewers, and no candidate is rejected by the system without human review.

3.2 Data, Provenance and Protection

The quantitative strand draws on the platform's operational logs and the organisations' HR systems. For each organisation we compared a three-month pre-deployment baseline with the six-month period after go-live, covering more than 6,000 requisitions. Three measures are used: time-to-hire (elapsed days from requisition approval to offer acceptance, which differs from time-to-fill); routine query volume (the number of policy or

administrative questions reaching human advisers, a measure of channel substitution rather than of underlying demand); and containment (the share of queries resolved by the platform without escalation, computed from interaction logs). Analytically, we summarise site-level reductions as medians and means with 95% confidence intervals (t-based, and read cautiously given only ten site-level observations), test the pre-post differences with the Wilcoxon signed-rank test, and report the matched-pairs rank-biserial correlation as an effect size; to confirm that no single organisation drives the pooled result, we recompute the medians under a leave-one-site-out (jackknife) procedure. The data are owned by the platform vendor and the participating organisations; the authors extracted and analysed them, harmonising differing HR-system definitions to common constructs. Query data were processed under each organisation's data-protection basis, then aggregated and de-identified before analysis. The figures derive from the vendor's own logs and were not independently audited, a limitation we return to.

3.3 Qualitative Strand

We conducted thirty-two semi-structured interviews (three to four per site) with practitioners in recruitment, HR operations and business-partner roles. Participants volunteered in response to an open invitation and were interviewed by researchers independent of their line management; participation was confidential. An interview guide covered changes in daily work, task movement, freed time and concerns. Transcripts were coded by two researchers using a combined inductive–deductive scheme; coding disagreements were reconciled, a subset was member-checked, and we actively sought disconfirming accounts. Themes are reported with frequencies (Table 3) but treated as self-reported perceptions subject to social-desirability pressure, particularly given the employer–vendor context.

3.4 Threats to Validity

Several threats temper interpretation. Sites and interviewees self-selected; adopters differ from non-adopters, and we are not aware of abandoned deployments in the sample, which risks survivorship bias. Gains may reflect concurrent process redesign rather than the platform alone. The mismatched three- and six-month windows expose the metrics to seasonality, and awareness of the study may have produced Hawthorne effects. Self-report and the vendor relationship (see Conflict of interest) may bias results in a favourable direction. We foreground these threats here, in the design, rather than confining them to a closing paragraph.

3.5 Governance of the Deployment

Recruitment AI of this kind is classified as high-risk under the EU AI Act (European Union 2024, Annex III), with obligations enforceable from August 2026, and is subject to Article 22 of the GDPR on solely-automated decisions (European Union 2016). In practice the deployment implemented several governance mechanisms: a human-in-the-loop review in which a named recruiter reviews every AI-generated shortlist and no candidate is rejected without human sign-off (the primary Article 22 safeguard); retention of decision logs and model-output audit trails to support the Act's logging and traceability duties; retrieval-augmented answers grounded in each organisation's own policy corpus to aid transparency and contestability; a documented data-protection basis and de-identification step for query data; and defined escalation paths for contested or low-confidence cases. Two obligations, however, were only partially met during the window. Systematic bias and adverse-impact testing of the screening tool, for example selection-rate or four-fifths-rule analysis across protected groups, was not performed uniformly across sites, and no formal fairness metric was tracked over time; nor was a full data-protection impact assessment completed everywhere. We report these gaps as findings and governance gaps, not as strengths, and return to them below as the central precondition for responsible scaling.

4. Findings

4.1 Operational Impact

Time-to-hire fell from a pooled mean of about 42 days before deployment to about 32 days afterwards, a median reduction of 23% across the ten sites (95% CI 20-26%; range 17–28%, SD 3.8), while routine queries reaching human advisers fell by a median of 72% (95% CI 69-75%; range 65–78%, SD 4.1) (Figure 1). A Wilcoxon signed-rank test across the ten sites indicated the pre–post differences were unlikely to be due to chance (n = 10 sites, exact two-sided p = 0.002), with every site moving in the same direction; the corresponding matched-pairs rank-biserial correlation is r = 1.00, which reflects the uniform direction of change rather than the size of any causal effect. With only ten site-level observations these inferential statistics are indicative rather than definitive, and the reductions proved robust in a leave-one-site-out check

(Section 4.2). The baseline is broadly in line with industry time-to-fill benchmarks of 36–44 days (SHRM 2025), though that is a different construct and the comparison is approximate. Table 1 summarises how task handling changed.

Table 1: HR tasks before and after platform deployment

HR task	Before deployment	After deployment	Headline change
CV screening and shortlisting	Manual, recruiter-led	AI-assisted; recruiter reviews shortlist	Time-to-hire −23% (median)
Interview scheduling	Manual coordination	Automated by the platform	Coordination time reduced
Policy and FAQ queries	Routed to human advisers	AI self-service; complex cases escalated	Routine queries to humans −72%
Onboarding, coaching, conflict, culture	Human, time-constrained	Human (unchanged in scope)	Self-reported capacity reallocated



Figure 1: Operational changes associated with the platform (indexed to pre-deployment baseline = 100; medians across ten sites)

4.2 Variation Across Sites

The direction of change was common to all ten sites, but its magnitude varied (Figure 2, Table 2): time-to-hire reductions ranged from 17% to 28%, query reductions from 65% to 78%, and containment from 65% to 83% (mean 74%, 95% CI 71-79%). A leave-one-site-out (jackknife) analysis confirmed the pooled result is not driven by any single organisation: the median time-to-hire reduction stayed within 22-24% and the median query reduction within 71-73% as each site was removed in turn. We therefore claim consistency of direction, not uniformity of magnitude, and we do not claim the sites are representative, since they are self-selected adopters. Per-site denominators are reported in Table 2; de-identified site-level data are available on reasonable request (see Data availability).

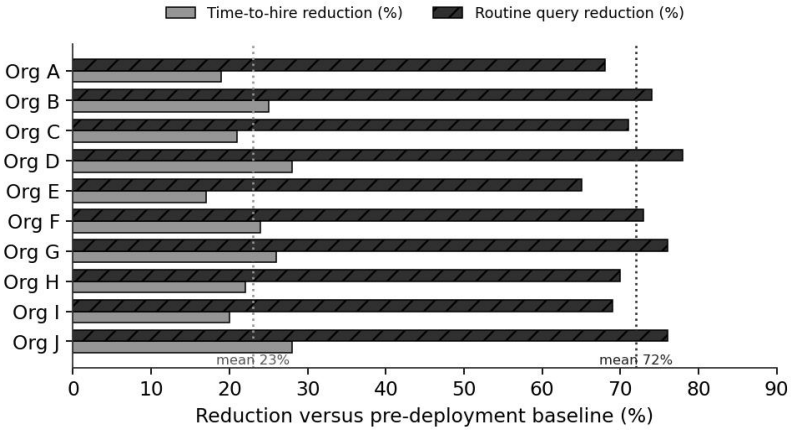


Figure 2: Reductions in time-to-hire and routine query volume by organisation (dotted lines show pooled means)

Table 2: Results by organisation (anonymised; n = hires processed in the window)

Org	Sector	Size	Hires (n)	TTH ↓	Query ↓	Contain.
A	Technology	5–20k	720	19%	68%	69%
B	Financial services	>20k	1,150	25%	74%	77%
C	Healthcare	5–20k	540	21%	71%	72%
D	Professional services	1–5k	210	28%	78%	83%
E	Public sector	>20k	980	17%	65%	65%
F	Retail	5–20k	610	24%	73%	75%
G	Telecommunications	>20k	1,040	26%	76%	79%
H	Manufacturing	5–20k	500	22%	70%	71%
I	Energy	1–5k	180	20%	69%	67%
J	Logistics	5–20k	470	28%	76%	78%
Mean / total	n/a	n/a	6,400	23%	72%	74%

4.3 Adoption Trajectory

The month-by-month trajectory is the study's most distinctive observation (Figure 3). Routine query volume fell fastest, dropping by about half within three months as employees learned to use self-service, then levelling off; time-to-hire declined more slowly and steadily as recruiters built confidence in AI-assisted shortlisting and adjusted their workflows. This lag is consistent with the idea that value follows the redesign of work rather than the moment of switch-on, but it is equally consistent with ordinary learning curves or selection effects, and the present design cannot distinguish these. The pattern nonetheless suggests that the human and process changes around the tool, not the tool alone, pace any benefits.

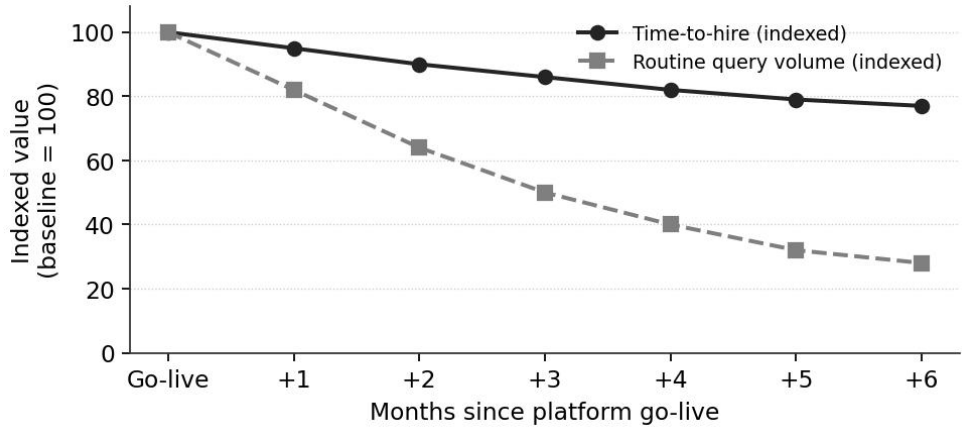


Figure 3: Time-to-hire and routine query volume over the first six months of use (indexed)

4.4 The Reconfigured Division of Labour

The platform did not absorb HR work evenly. It took over high-volume, rule-bound tasks such as screening and shortlisting, scheduling, request routing and standardised policy answers, while tasks depending on judgement, trust and relationship remained human: onboarding, coaching, conflict resolution and the work of sustaining culture. In Acemoglu and Restrepo's (2019) terms, displacement fell on routine tasks while the reinforced tasks are those in which labour retains an advantage. Practitioners' self-reported time allocation shifted accordingly (Figure 4): in a time-allocation question put to the thirty-two interviewees, estimated time on transactional work fell from about 55% to 28%, while relational work rose from 18% to 35% and strategic or advisory work increased from 7% to 15%. We define transactional work as rule-based administration, relational work as direct people support, and strategic work as planning and advisory activity. These are self-reported estimates from a small sample, not time-and-motion measurements; the six-month swing is large and should be read with caution.

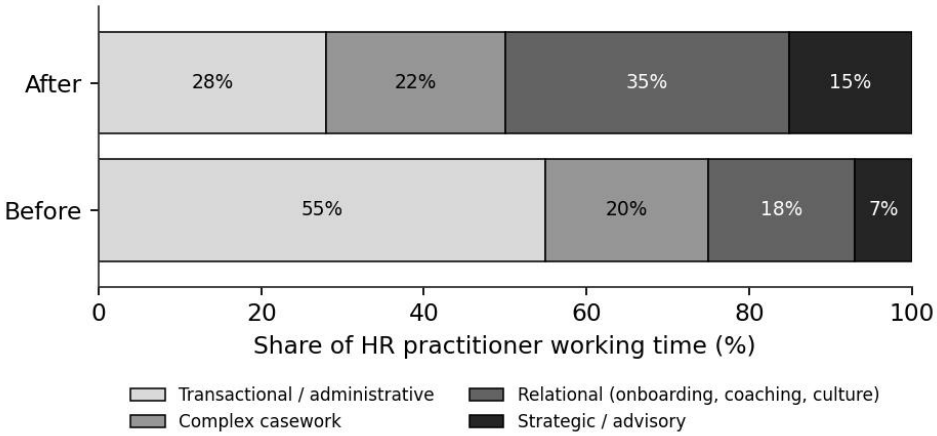


Figure 4: Self-reported reallocation of HR practitioner time before and after deployment (n = 32)

4.5 Practitioner Experience

Six themes recurred (Table 3). Practitioners most often described regaining time for engagement and a shift towards complex, judgement-based work, and many reported needing new skills to supervise and correct the system. The accounts were not uniformly positive. A distinct minority raised concerns consistent with the critical literature on algorithmic management (Kellogg, Valentine and Christin 2020): that the system could be used to monitor how quickly staff close cases, that expectations simply rose to fill the freed time, and that some screening craft was being lost. An initial fear of redundancy featured in several interviews; it eased for many as roles expanded but did not disappear, and we do not read it as resolved. These are perceptions, and the employer–vendor context makes social-desirability bias a real risk; we report the disconfirming accounts precisely because the setting predisposes towards favourable ones. The positive themes align with work linking employee–AI collaboration to engagement when work is redesigned (Sun, Hu and Su 2026; Watermann, Kubowitsch and Lermer 2025), while the critical accounts caution against treating that outcome as inevitable (Viljakainen et al 2025).

Table 3: Themes from interviews with HR practitioners (n = 32)

Theme	Count	Illustrative quotation
Time released for candidate and employee engagement	21	"I finally have time to talk to candidates instead of chasing paperwork."
Shift towards complex, judgement-based work	18	"The cases that reach me now are the ones that genuinely need a person."
New supervisory and AI-related skills required	15	"We had to learn to check and correct the system, and that is a new part of the job."
Transparency as a condition for trust	14	"Once we explained how the assistant reached its answers, people trusted it."
Initial fear of redundancy, eased but not removed	12	"People feared it would replace us; for many that eased as roles grew, but the worry never fully went."
Concern about surveillance, intensification or deskilling	8	"There's unease it tracks how fast we close cases, and that some screening craft is being lost."

5. Discussion

Within these ten organisations, the platform was associated with a redistribution of HR work: routine tasks moved to the machine while relational tasks remained and, in self-report, expanded. Read through the task-based lens this resembles reinstatement keeping pace with displacement (Acemoglu and Restrepo 2019), and through AMO and work-design theory it resembles enrichment of HR roles (Appelbaum et al 2000; Bakker and Demerouti 2007). But these are six-month associations from self-selected sites, not evidence that augmentation is the lasting equilibrium; automation and augmentation remain interdependent (Raisch and Krakowski 2021). To be explicit about inference: because the design has no control group or counterfactual, these results identify associations between deployment and subsequent change, not causal effects; concurrent process redesign, ordinary learning curves, seasonality, and selection into adoption remain equally plausible

explanations, and we use the language of movement and redistribution descriptively rather than causally throughout.

Theoretical contribution. Beyond applying these lenses, we intend a specific contribution to theory. The automation-augmentation paradox has largely been argued conceptually or at the level of the individual task (Raisch and Krakowski 2021); our multi-site evidence adjudicates it at the organisational level, showing displacement of routine tasks and reinstatement of relational ones co-occurring within the same function and time window rather than as opposing trajectories. In doing so we make explicit a boundary condition that the task-based account (Acemoglu and Restrepo 2019) leaves open: reinstatement did not follow automatically from deployment but tracked complementary work redesign, which connects the sociotechnical premise (Trist and Bamforth 1951) to the empirical '10-20-70' pattern and reframes augmentation as a contingent, design-dependent outcome rather than a property of the technology. This is a scope condition and an organisational-level test of existing theory rather than a wholly new construct, and we state it as such.

The gradual trajectory speaks to why so many AI programmes show little measurable impact (MIT NANDA 2025). Industry analyses of AI transformation (the '10-20-70' rule) hold that around 70% of the effort needed to capture value lies in people and process change rather than in the algorithm or infrastructure (BCG 2026). Our data are consistent with this as a hypothesis: benefits tracked the redesign of work, not the switch-on. We cannot demonstrate the link, however, because we did not compare redesigned with non-redesigned sites.

For a governance audience the implications are concrete. Because recruitment AI is high-risk under the EU AI Act (European Union 2024) and engages Article 22 of the GDPR (European Union 2016), deployments require documented human oversight, logging and, critically, bias and adverse-impact testing of screening tools. Practically this means routine measurement of selection rates and outcomes across protected groups (for example against the four-fifths rule), pre-deployment and periodic fairness audits, model and data documentation, and a completed data-protection impact assessment, with the results fed back into model configuration. Our sites relied on human-in-the-loop review and logging but had not systematically implemented this fairness testing; closing that gap is a precondition for responsible scaling, not an optional extra (Taherdoost 2025; Sony et al 2025; Alameer 2024).

Finally, the claim that efficiency can fund new, AI-enabled roles rather than redundancies is a recommendation, not a finding: we measured neither new roles created nor reskilling budgets, and macro projections show large gross job losses alongside larger gross gains, with a contested net (WEF 2025; Goldman Sachs 2025). What our study supports is narrower: that freed capacity exists and, within the window, was redirected towards higher-value work. Whether leaders convert it into capability or into headcount reduction is a strategic and governance choice this evidence cannot make for them.

A final limitation warrants critical reflection rather than mere disclosure. The platform is a product of the authors' employer, the participating organisations are its customers, and the operational figures originate from vendor-owned systems that were not independently audited. Each of these could bias the findings in a favourable direction, through identifiable mechanisms: selection, because organisations that adopted and persisted with the platform are unlikely to represent those that trialled and abandoned it; measurement, because vendor-defined log constructs may flatter the tool and could not be externally verified; and social desirability, because interviewees are employees of customers speaking about their supplier. We have tried to counter these tendencies by scoping claims to the data, reporting associations rather than effects, foregrounding threats to validity in the design, actively seeking disconfirming interview accounts, and offering de-identified site-level data for inspection. These mitigations reduce but cannot remove the risk; accordingly the results should be read as vendor-context, transparency-qualified evidence, and independent, third-party replication and audit are needed before the pattern is generalised.

6. Conclusion

We asked how generative AI redistributes HR work. In an exploratory, pre–post study of one platform across ten organisations over six months, deployment was associated with a 23% median fall in time-to-hire and a 72% fall in routine queries reaching advisers, accompanied by the automation of routine tasks and the reinforcement of relational ones; practitioners reported redirecting freed time to higher-value work, with notable concerns about surveillance, intensification and trust. The contribution is a rare, if early, organization-level account of task redistribution. Its limits are intrinsic: a single vendor-connected platform, self-selected sites, self-report, a six-month window too short to observe employment outcomes, and operational data that were not independently audited. We therefore frame 'augmentation over attrition' not as a result but as a

conditional choice that depends on leadership, governance and reskilling. Future work should use longitudinal, multi-platform and ideally controlled designs, compare redesigned with non-redesigned sites, incorporate applicant perspectives, and independently replicate and audit deployed systems for bias.

7. Conflict of Interest, Funding, and Data

The enterprise generative AI platform evaluated here is a commercial product developed by the authors' employer, and several of the ten participating organizations are clients of that employer. The research was conducted as part of the authors' employment and received no external funding. The operational data are owned by the vendor and the participating organizations; the authors extracted and analyzed them. The findings are commercially favorable to the platform, and readers should weigh them against this relationship; we have sought to mitigate it by scoping claims to the data, reporting results as associations, foregrounding limitations, and offering data for inspection. We accordingly encourage independent, third-party replication and audit of both the operational figures and the screening tool, and we present the findings as vendor-context evidence to be weighed on those terms.

8. Data Availability and Independence

De-identified, aggregated per-site pre/post values and denominators are available from the authors on reasonable request. The operational figures derive from the platform's own logs and were not subject to an independent third-party audit; this constrains their evidential weight and is stated as a limitation.

Ethics Declaration: Participation in interviews was voluntary and based on written informed consent; interviews were conducted by researchers independent of participants' line management, responses were anonymized, and no identifiable data were shared with employers, in order to manage the employee–employer–vendor power dynamic. Operational data were aggregated and de-identified.

AI Declaration: Generative AI tools such as Claude and ChatGPT were used to support literature searching, structuring, and language editing in preparing this paper; all AI-generated text was reviewed, verified, and edited by the authors, who are responsible for all content. Separately, and distinct from this writing assistance, the study analyses data from a commercial AI platform connected to the authors; that relationship and the ownership of the operational data are disclosed above under Conflict of interest.

References

Abdelhay, S., AlTalay, M.S.R., Selim, N., Altamimi, A.A., Hassan, D., Elbannany, M. and Marie, A. (2025) "The impact of generative AI (ChatGPT) on recruitment efficiency and candidate quality: the mediating role of process automation level and the moderating role of organizational size", *Frontiers in Human Dynamics*, Vol 6, article 1487671. doi:10.3389/fhumd.2024.1487671.

Acemoglu, D. and Restrepo, P. (2019) "Automation and new tasks: how technology displaces and reinstates labor", *Journal of Economic Perspectives*, Vol 33, No. 2, pp 3–30. doi:10.1257/jep.33.2.3.

Alameer, A. (2024) "A comprehensive review of AI techniques for addressing algorithmic bias in job hiring", *AI*, Vol 5, No. 1, pp 383–404. doi:10.3390/ai5010019.

Appelbaum, E., Bailey, T., Berg, P. and Kalleberg, A.L. (2000) *Manufacturing Advantage: Why High-Performance Work Systems Pay Off*, Cornell University Press, Ithaca, NY.

Bakker, A.B. and Demerouti, E. (2007) "The job demands-resources model: state of the art", *Journal of Managerial Psychology*, Vol 22, No. 3, pp 309–328. doi:10.1108/02683940710733115.

Boston Consulting Group (2026) *AI Transformation Is a Workforce Transformation*, BCG, Boston. Available at https://www.bcg.com/publications/2026/ai-transformation-is-a-workforce-transformation (Accessed 26 June 2026).

Cedro, T.L. (2025) "The ethical role of generative artificial intelligence in modern HR decision-making: a systematic literature review", *European Journal of Business and Management Research*, Vol 10, No. 1, pp 44–55. doi:10.24018/ejbmr.2025.10.1.2535.

Daugherty, P.R. and Wilson, H.J. (2018) *Human + Machine: Reimagining Work in the Age of AI*, Harvard Business Review Press, Boston, MA.

Davenport, T.H. and Kirby, J. (2016) *Only Humans Need Apply: Winners and Losers in the Age of Smart Machines*, Harper Business, New York.

European Union (2016) *Regulation (EU) 2016/679 (General Data Protection Regulation)*, Official Journal of the European Union, L 119, pp 1–88. Available at https://eur-lex.europa.eu/eli/reg/2016/679/oj (Accessed 26 June 2026).

European Union (2024) *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*, Official Journal of the European Union, L 2024/1689. Available at https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689 (Accessed 26 June 2026).

Goldman Sachs (2025) *How Will AI Affect the Global Workforce?*, Goldman Sachs, New York. Available at https://www.goldmansachs.com/insights/articles/how-will-ai-affect-the-global-workforce (Accessed 26 June 2026).

IBM (2025) *IBM AskHR*, IBM Case Studies, Armonk, NY. Available at https://www.ibm.com/case-studies/ibm-askhr (Accessed 26 June 2026).

International Labour Organization (2025) *Generative AI and Jobs: A 2025 Update*, ILO, Geneva. Available at https://www.ilo.org/publications/generative-ai-and-jobs-2025-update (Accessed 26 June 2026).

Kellogg, K.C., Valentine, M.A. and Christin, A. (2020) "Algorithms at work: the new contested terrain of control", *Academy of Management Annals*, Vol 14, No. 1, pp 366–410. doi:10.5465/annals.2018.0174.

Madanchian, M. (2024) "From recruitment to retention: AI tools for human resource decision-making", *Applied Sciences*, Vol 14, No. 24, article 11750. doi:10.3390/app142411750.

Mamadjanova, K. (2025) "Role of artificial intelligence in employee recruitment: systematic review and future research directions", *Discover Global Society*, Vol 3, No. 1, article 99. doi:10.1007/s44282-025-00246-w.

MIT NANDA (2025) *The State of AI in Business 2025*, Massachusetts Institute of Technology, Cambridge, MA. Available at https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf (Accessed 26 June 2026).

Raisch, S. and Krakowski, S. (2021) "Artificial intelligence and management: the automation–augmentation paradox", *Academy of Management Review*, Vol 46, No. 1, pp 192–210. doi:10.5465/amr.2018.0072.

Rukadikar, A., Khandelwal, K. and Warrier, U. (2025) "Reimagining recruitment: traditional methods meet AI interventions – a 20-year assessment (2003–2023)", *Cogent Business & Management*, Vol 12, No. 1, article 2454319. doi:10.1080/23311975.2025.2454319.

SHRM (2025) *2025 Recruiting Benchmarking Report*, Society for Human Resource Management, Alexandria, VA. Available at https://www.shrm.org/topics-tools/research (Accessed 26 June 2026).

Sony, M.M.A.A.M., Amin, M.B., Ashraf, A., Islam, K.M.A., Debnath, N.C. and Debnath, G.C. (2025) "Bias in AI-driven HRM systems: investigating discrimination risks embedded in AI recruitment tools and HR analytics", *Social Sciences & Humanities Open*, Vol 12, article 102082. doi:10.1016/j.ssaho.2025.102082.

Sun, L., Hu, R. and Su, H. (2026) "Unlocking human potential in the AI age: how employee-AI collaboration transforms work engagement through dual psychological pathways", *Frontiers in Psychology*, Vol 16, article 1705671. doi:10.3389/fpsyg.2025.1705671.

Taherdoost, H. (2025) "Ethical theories, governance models, and strategic frameworks for responsible AI adoption and organizational success", *Frontiers in Artificial Intelligence*, Vol 8, article 1619029. doi:10.3389/frai.2025.1619029.

Tambe, P., Cappelli, P. and Yakubovich, V. (2019) "Artificial intelligence in human resources management: challenges and a path forward", *California Management Review*, Vol 61, No. 4, pp 15–42. doi:10.1177/0008125619867910.

Trist, E.L. and Bamforth, K.W. (1951) "Some social and psychological consequences of the longwall method of coal-getting", *Human Relations*, Vol 4, No. 1, pp 3–38.

Venkatesh, V., Morris, M.G., Davis, G.B. and Davis, F.D. (2003) "User acceptance of information technology: toward a unified view", *MIS Quarterly*, Vol 27, No. 3, pp 425–478. doi:10.2307/30036540.

Venugopal, M., Madhavan, V., Prasad, R. and Raman, R. (2024) "Transformative AI in human resource management: enhancing workforce planning with topic modeling", *Cogent Business & Management*, Vol 11, No. 1, article 2432550. doi:10.1080/23311975.2024.2432550.

Viljakainen, A., Silmukari, J., Roto, V., Čaić, M. and Aromaa, S. (2025) "AI and the future of industrial work: a framework for enhancing employee experience from satisfaction to flourishing", *Behaviour & Information Technology*, pp 1–13. doi:10.1080/0144929X.2025.2596889.

Vrontis, D., Christofi, M., Pereira, V., Tarba, S., Makrides, A. and Trichina, E. (2022) "Artificial intelligence, robotics, advanced technologies and human resource management: a systematic review", *The International Journal of Human Resource Management*, Vol 33, No. 6, pp 1237–1266. doi:10.1080/09585192.2020.1871398.

Watermann, L., Kubowitsch, S. and Lermer, E. (2025) "AI and work design: a positive psychology approach to employee well-being", *Gruppe. Interaktion. Organisation. Zeitschrift für Angewandte Organisationspsychologie (GIO)*, Vol 56, No. 2, pp 311–320. doi:10.1007/s11612-025-00806-3.

World Economic Forum (2025) *The Future of Jobs Report 2025*, World Economic Forum, Geneva. Available at https://www.weforum.org/publications/the-future-of-jobs-report-2025/ (Accessed 26 June 2026).