



(RESEARCH ARTICLE)



Navigating the AI revolution in talent acquisition: A qualitative, hypothesis-driven study of adoption drivers, regulatory barriers and governance imperatives across global organizations

Nnanna John *

Engineering Science and Technology Entrepreneurship, University of Notre Dame.

World Journal of Advanced Research and Reviews, 2026, 30(02), 048-072

Publication history: Received on 25 March 2026; revised on 30 April 2026; accepted on 02 May 2026

Article DOI: <https://doi.org/10.30574/wjarr.2026.30.2.1165>

Abstract

Artificial intelligence (AI) is transforming talent acquisition (TA) practices across global organizations, yet empirical evidence on the organizational, regulatory, and governance dimensions of this transformation remains fragmented. Against a backdrop of global RPO market restructuring (Everest Group, 2024) used here as contextual market intelligence, this study presents findings from a qualitative, hypothesis-driven interview investigation involving 50 in-depth stakeholder interviews spanning 14 industries across North America, Europe, and Asia. Data were captured through video recordings and transcripts; AI-assisted transcript analysis was employed to surface themes across the dataset. A hypothesis-driven protocol grounded in the Technology Acceptance Model (TAM) and Institutional Theory systematically validated and invalidated 15 propositions. Key findings reveal that AI integration in candidate screening reduces time-to-fill (validated across 39 interviews); regulatory concerns, particularly the EU AI Act (Regulation 2024/1689) and New York City's Automated Employment Decision Tool (AEDT) regulation constitute the primary adoption barrier in regulated industries (29 interviews); executive governance structures are a decisive implementation success factor; and skills-based hiring is gaining traction as an AI-enabled paradigm shift. Three hypotheses were invalidated, including premium pricing for speed and inclusion features and the belief that unbundled tools drive higher adoption. These findings contribute a theoretically grounded, cross-industry empirical perspective to the literature on responsible AI deployment in talent acquisition.

Keywords: Artificial Intelligence; Talent Acquisition; Recruitment Process Outsourcing; Technology Acceptance Model; EU AI Act; AEDT; Skills-Based Hiring; AI- Retrieval; HR Governance

1. Introduction

The integration of artificial intelligence into organizational talent acquisition processes represents one of the most consequential technological shifts in human resource management in recent decades. In the World Economic Forum's Future of Jobs Report 2025, AI and big data skills are ranked among the fastest-growing competencies, and 39% of workers' core skills are projected to change by 2030 (World Economic Forum, 2025). The report, drawing on survey data from over 1,000 employers representing more than 14 million workers across 55 economies projects that 170 million new jobs will be created globally by 2030 while 92 million roles will be displaced (World Economic Forum, 2025). Employers project that 63% identify skills gaps as the primary barrier to business transformation. For talent acquisition professionals, these are not distant forecasts; they are present operational imperatives.

Despite this urgency, the RPO industry providing much of global enterprise hiring infrastructure has experienced significant turbulence. Everest Group's State of the Market 2024 documents that following post-pandemic growth, the global RPO market declined sharply, driven by macroeconomic uncertainty, geopolitical tensions, and shifting

* Corresponding author: Nnanna John

enterprise priorities - a slowdown that continued into 2024, prompting widespread workforce strategy realignments (Everest Group, 2024). In this context, AI has emerged not merely as an efficiency tool but as a strategic differentiator: the Everest Group PEAK Matrix Assessment 2025 examining 31 providers across more than 6,700 multi-process RPO deals identifies technology integration and innovation investment as primary differentiators among global market leaders (Everest Group, 2025).

The regulatory environment compounds this complexity. The EU AI Act (Regulation 2024/1689), which entered into force in August 2024, classifies employment AI as 'high-risk' and introduces conformity assessments, transparency requirements, and human oversight mandates. Its extraterritorial reach applies to any organization deploying AI in hiring decisions involving EU residents (European Commission, 2024). New York City's Local Law 144 (2021) regulates the use of Automated Employment Decision Tools (AEDTs) in hiring by requiring employers to conduct annual independent bias audits, publicly disclose audit results, and notify candidates at least ten business days prior to AEDT use. Non-compliance is enforced by the NYC Department of Consumer and Worker Protection (DCWP), with civil penalties beginning at \$500 for an initial violation and escalating to \$1,500 per day for subsequent violations.

This study responds to the convergence of these market, regulatory, and organizational dynamics. Drawing on 50 structured interviews across 14 industries in North America, Europe, and Asia, with data captured via video recordings and transcripts and analyzed using AI-assisted transcript review, the study employs a qualitative, hypothesis-driven design to examine AI adoption in TA through theory-derived propositions. The study makes three primary contributions: it provides a substantive cross-industry primary interview dataset in this domain; it situates findings alongside widely cited practitioner-facing market analyses (e.g., Everest Group) as contextual industry intelligence rather than independent corroboration; and it systematically tests and invalidates three widely held assumptions, producing corrective guidance for practitioners, vendors, and investors.

The paper proceeds as follows. The literature review is presented in Section 2, followed by the theoretical framework in Section 3 and the methodology in Section 4. Section 5 reports the findings, Section 6 discusses implications, and Section 7 concludes.

2. Literature Review

2.1. The Macro-Context: AI, Skills Disruption, and the Future of Work

The World Economic Forum's Future of Jobs Report 2025 provides the most authoritative current mapping of labour market transformation relevant to talent acquisition. Drawing on survey data from over 1,000 employers representing more than 14 million workers across 22 industry clusters and 55 economies, the report identifies five macro-trends reshaping labour markets by 2030: technological change, the green transition, geoeconomic fragmentation, economic uncertainty, and demographic shifts (WEF, 2025). AI and big data top the list of fastest-growing core skills. Employers project that 39% of workers' core skills will change by 2030 down from 44% projected in 2023, reflecting growing investment in upskilling and 85% of employers plan upskilling as their primary workforce strategy. Skill gaps are the leading transformation barrier for 63% of employers, and only 29% of businesses anticipate improved talent availability between 2025 and 2030 (WEF, 2025). However, because the WEF evidence is based on employer surveys and macro-level aggregation, it offers limited visibility into how talent acquisition leaders operationalize AI under real workflow, governance, and regulatory constraints, an empirical gap addressed by this study's multi-stakeholder interviews.

These figures establish the structural urgency of this study's research question: as skills requirements shift rapidly, the talent acquisition function that identifies and assesses those skills must evolve in parallel. Critically, the WEF report emphasizes human-machine collaboration, not displacement as the dominant future model: every industry is projected to see a decrease in tasks performed exclusively by humans, but augmentation rather than wholesale automation is the primary mechanism. This framing aligns closely with the consensus across this study's 50 interviews, as reported in the findings. At the same time, the interviews refine the augmentation claim by showing that augmentation in TA is uneven and conditional, often limited to administrative tasks, slowed by regulatory and liability concerns in determinative decision points, and dependent on governance and training structures that survey-based macro accounts cannot readily observe.

2.2. The RPO Market: Disruption, Resilience, and AI as Strategic Differentiator

Everest Group's State of the Market 2024, analyzing over 6,400 active multi-process RPO deals as of 2023 documents a sector in structural adjustment. Following post-pandemic growth, the global RPO market declined sharply. In response, providers pivoted toward greater agility, investing in modular solutions, consulting and advisory services, and critically,

skills-based hiring capabilities as a primary buyer expectation (Everest Group, 2024). The report identifies the growing prominence of skills-based hiring and the rise of strategic workforce advisory as defining shifts in how RPO providers position their value.

The Everest Group PEAK Matrix Assessment 2025 confirms this strategic repositioning. Examining 31 providers across more than 6,700 multi-process RPO deals as of 2024, the assessment identifies technology integration, innovation investment, and AI-enabled capabilities as primary differentiators among market leaders. Providers that invested in AI-powered screening, generative AI for candidate engagement, and skills-based recruitment architectures maintained competitive positioning during a period of market-wide demand contraction (Everest Group, 2025). This market evidence establishes that AI adoption in TA is no longer experimental; it is competitive.

2.3. AI Adoption Patterns in HR

The academic literature on AI adoption in HR has expanded significantly in recent years, though it remains geographically concentrated and predominantly focused on adoption intentions rather than organizational outcomes. A recent study examining effective AI adoption in talent acquisition across multinational corporation contexts finds that adoption is shaped by interacting individual factors, including fear of workflow complexity and bias concerns and organizational factors including competitive pressure and regulatory environment, arguing for integration of TAM with the Technology-Organization-Environment (TOE) framework (Roppelt et al., 2025). However, this body of work largely emphasizes determinants of adoption rather than testing whether adoption produces validated performance outcomes in practice; this study extends the literature by examining outcome claims (e.g., efficiency gains and their limits) across 50 interviews spanning multiple industries and three continents.

The foundational challenge framework for AI in HRM was established by Tambe et al. (2019), who identify four structural barriers to realizing AI's promise in HR contexts: the complexity of HR phenomena, small and biased data sets, accountability and fairness constraints, and adverse employee reactions to algorithmic decision-making. These barriers remain structurally relevant to the talent acquisition context examined in this study. Yet prior work typically articulates these barriers at a conceptual level rather than tracing how they translate into concrete executive decisions under contemporary regulatory regimes (e.g., the EU AI Act and AEDT) and internal governance requirements, an implementation-level gap this study addresses through interview evidence.

A qualitative study employing grounded theory to examine AI's organizational and operational dimensions in talent acquisition conducted with HR professionals and AI platform providers in Sweden identifies four aggregate dimensions of AI's transformative role: operational efficiency, candidate experience, decision quality, and organizational learning (Paramita et al., 2024). Their findings are consistent with the augmentation framing evident across this study's interview data. However, single-country qualitative designs cannot fully capture the cross-jurisdictional compliance pressures and competitive market dynamics that shape AI deployment decisions across regions; by sampling across North America, Europe, and Asia and including executives and investors, this study extends the context and unit-of-analysis coverage of prior qualitative work.

Early literature projected a broadly optimistic trajectory for AI in recruitment, with Hmoud and Várallyai (2019) concluding that AI provides promising solutions for recruiters by automating time-consuming repetitive tasks such as sourcing and screening, improving hiring quality and reducing human bias. This optimism, however, is conditioned on intentional governance, a point taken up directly by research on inclusive hiring and bias, which argues that inclusivity in AI is not automatic but requires deliberate design and multi-stakeholder collaboration (Ezenwa, 2025), without which AI risks encoding and amplifying the historical biases present in training data. What remains underexplored is how these optimistic claims intersect with market and implementation realities, particularly whether organizations will pay for inclusion and speed features, and what negative downstream consequences emerge when AI is deployed without sufficient safeguards, questions this study addresses through hypothesis testing and cross-stakeholder interview evidence.

2.4. Regulatory Dimensions: EU AI Act, AEDT, and GDPR

The regulatory landscape governing AI in employment has formalized rapidly. The EU AI Act (Regulation 2024/1689), entering into force August 2024, classifies employment-related AI systems as 'high-risk,' triggering conformity assessments, transparency requirements, human oversight mandates, and registration obligations. Prohibited practices, including emotion recognition in workplace and hiring contexts became effective February 2025. Core high-risk system requirements phase in from August 2026, giving organizations a defined compliance runway. The Act's extraterritorial reach means organizations anywhere in the world using AI in hiring decisions involving EU residents are subject to its provisions (European Commission, 2024).

New York City's Local Law 144 (2021) regulates the use of Automated Employment Decision Tools (AEDTs) in hiring by requiring employers to conduct annual independent bias audits, publicly disclose audit results, and notify candidates at least ten business days prior to AEDT use. Non-compliance is enforced by the NYC Department of Consumer and Worker Protection (DCWP), with civil penalties beginning at \$500 for an initial violation and escalating to \$1,500 per day for subsequent violations. The European Parliament Research Service has flagged unresolved legal tension between the AI Act's provisions allowing processing of special categories of personal data for bias detection and the GDPR's more restrictive approach to such processing (De Luca & Federico, 2025).

2.5. Adoption Theory: TAM and Institutional Theory

The Technology Acceptance Model (Davis, 1989) posits that perceived usefulness and perceived ease of use are the primary determinants of technology adoption intentions. Applied to AI in TA, TAM predicts that recruiters and HR leaders adopt AI tools when they believe those tools will improve hiring outcomes and when the tools are not perceived as prohibitively complex. The study's training-related finding that recruiter training investment directly predicts adoption outcomes maps precisely onto TAM's perceived ease of use construct: training converts a tool from perceived complexity to perceived accessibility.

Institutional Theory (DiMaggio & Powell, 1983) provides the organizational complement, explaining how three distinct isomorphic pressures produce divergent AI adoption responses even among organizations in equivalent competitive positions. Three isomorphic pressures are relevant: coercive isomorphism, in which regulatory mandates - the EU AI Act, AEDT - shape adoption behavior regardless of individual-level usefulness calculations; mimetic isomorphism, in which competitive pressure drives adoption as organizations emulate peers; and normative isomorphism, in which professional standards and training norms establish AI fluency as an expected practitioner competency.

2.6. Skills-Based Hiring as an AI-Enabled Paradigm Shift

Skills-based hiring represents a structural departure from credential-based and job-title-based evaluation frameworks that have dominated recruitment practice for decades. Rather than using educational qualifications or previous job titles as proxies for capability, skills-based approaches evaluate candidates on demonstrated competencies directly relevant to role performance. This shift has theoretical grounding in human capital theory (Becker, 1964), which holds that productive capacity derives from specific skills rather than credentialed status. It also connects to the broader literature on person-job fit, which identifies competency alignment, rather than demographic or credential similarity, as the strongest predictor of job performance and retention. The operationalization of skills-based hiring at scale has historically been constrained by the difficulty of assessing skills consistently across large candidate pools; AI-powered skills matching and ontology frameworks are now removing that constraint, making the theoretical ideal practically achievable for the first time. That this shift is occurring in practice, not just in theory, is confirmed by the market evidence reviewed in Section 2.2.

2.7. Research Gaps

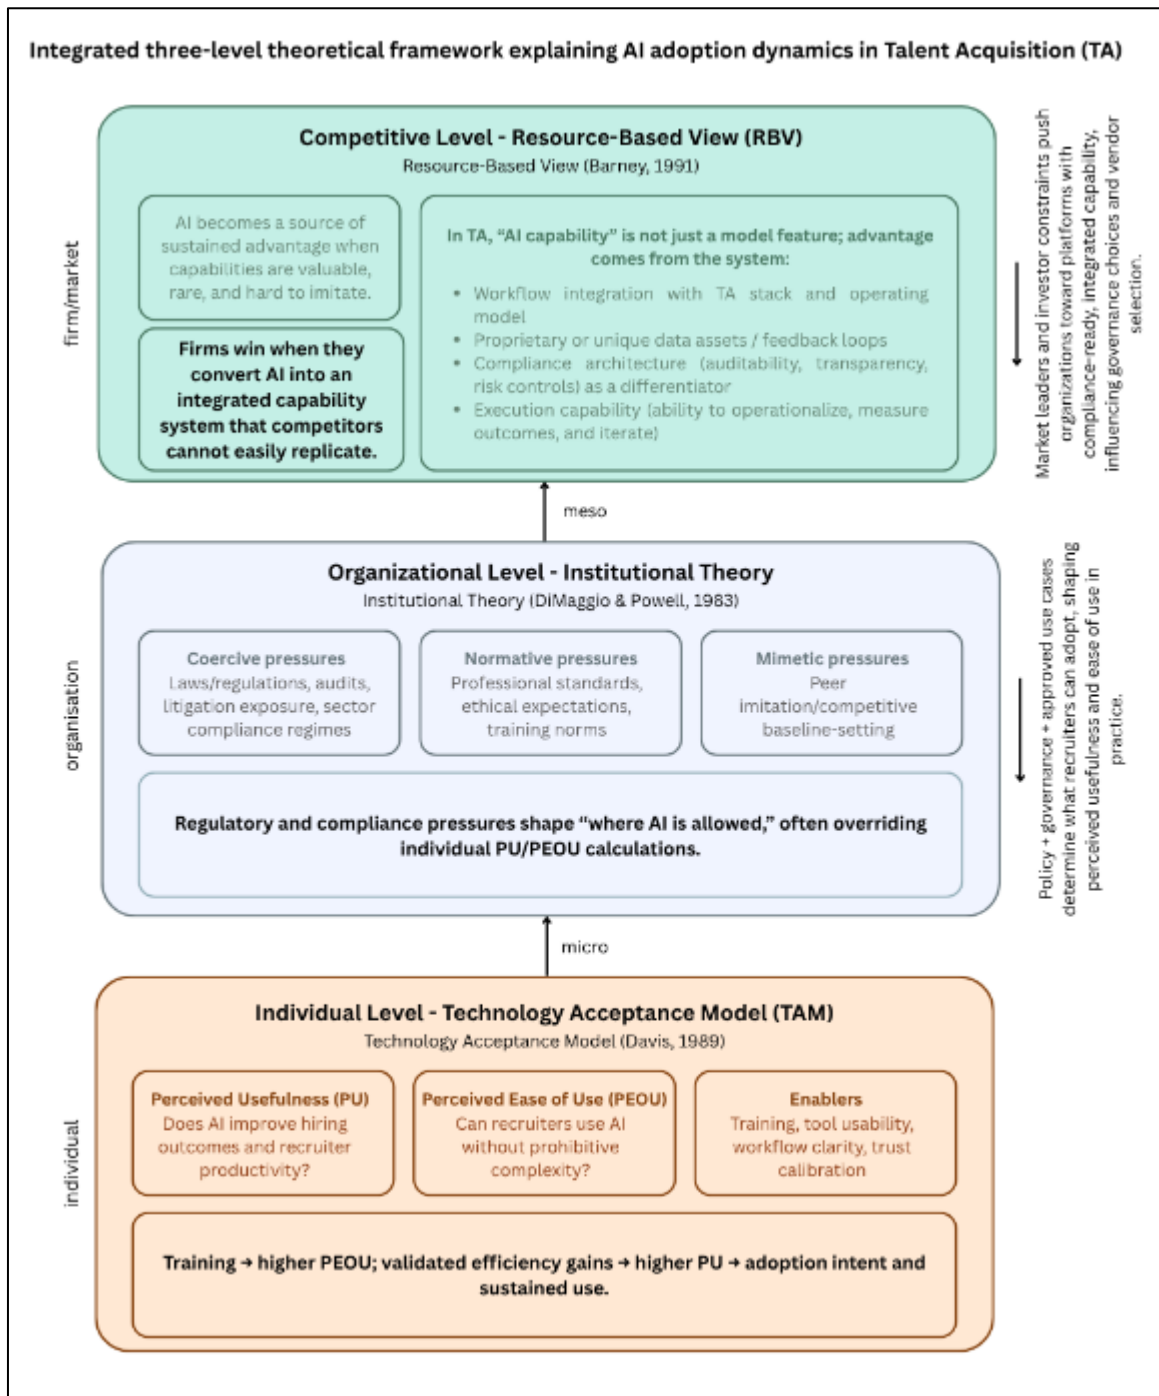
Three gaps motivate this study. First, most existing empirical research is geographically concentrated or sector-specific, limiting generalizability across industries and regions. Second, the literature focuses predominantly on adoption intentions rather than validated organizational outcomes, this study tests 15 hypotheses against actual interview evidence. Third, while systematic reviews of AI in recruitment and selection document a rapidly growing body of work focused on practitioner adoption, candidate perceptions, ethical risk, and regulatory questions, they rarely examine how investors evaluate AI-enabled TA companies or how investment criteria shape the TA technology ecosystem (Mori et al., 2025; Herold & Roedenbeck, 2025). This study addresses all three gaps.

3. Theoretical Framework

This study employs an integrated theoretical framework combining the Technology Acceptance Model (Davis, 1989) at the individual level, Institutional Theory (DiMaggio & Powell, 1983) at the organizational level, and the Resource-Based View (Barney, 1991) at the competitive level. Each operates at a different unit of analysis; together they account for the full range of dynamics observed across the 50 interviews.

At the individual level, TAM explains recruiter and HR leader adoption in terms of perceived usefulness, does AI meaningfully improve time-to-fill, candidate quality, or efficiency? and perceived ease of use, can the tool be used without prohibitive complexity? The study's training finding maps directly onto the perceived ease of use construct: training lowers the perceived complexity of AI tools, thereby removing a primary adoption barrier.

At the organizational level, Institutional Theory explains why organizations in equivalent competitive positions respond differently to AI. Coercive isomorphism, through regulatory mandates from the EU AI Act and AEDT produces adoption conservatism in regulated industries, irrespective of individual-level usefulness calculations. Mimetic isomorphism - competitive pressure - produces adoption behavior driven by peer imitation rather than internal ROI analysis, consistent with the Everest Group (2025) finding that AI investment is becoming a market baseline among top-tier providers. Normative isomorphism - professional standards and training norms - explains the role of recruiter training programs in establishing AI fluency as an expected competency.



Note: The framework integrates Technology Acceptance Model (TAM), Institutional Theory, and the Resource-Based View (RBV) to explain AI adoption in talent acquisition at three units of analysis. Section 5 (Individual: 5.1; Organizational: 5.2; Competitive: 5.4) provides validated findings.

Figure 1 Integrated three-level theoretical framework explaining AI adoption dynamics in Talent Acquisition (TA)

At the competitive level, the Resource-Based View (Barney, 1991) frames AI capability as a source of sustained competitive advantage when it is valuable, rare, and difficult to imitate. The Everest Group PEAK Matrix findings are

consistent with this framing: leading RPO providers are distinguished by proprietary AI capabilities that clients cannot easily replicate. The invalidation of the unbundling hypothesis also has RBV implications: buyers prefer integrated solutions because the competitive value of AI in TA is realized through systemic integration with existing workflows, not through standalone point-tool deployment.

4. Research Methodology

4.1. Research Design

This study employs a qualitative, hypothesis-driven research design based on deductive qualitative analysis (Gilgun, 2019). The interview guide operationalized theory-derived propositions and assessed the extent to which interview evidence supported, refined, or contradicted those propositions across the sample. To increase transparency regarding evidentiary strength, the analysis also reports descriptive support levels (i.e., the number of interviews providing relevant evidence for each proposition). These counts are used as within-method descriptors to summarize how widely a pattern was observed across stakeholder interviews, rather than as a separate quantitative strand or statistical inference.

4.2. Sampling Strategy

Purposive sampling was employed to maximize variation across stakeholder type, industry, organizational size, and geography (Patton, 2002). Four stakeholder categories were represented: (1) talent acquisition leaders, corporate recruiters, and HR executives from global enterprises spanning Information Technology, Financial Services, Healthcare, Staffing, Construction, Higher Education, Marketing, and other sectors; (2) Chief Human Resource Officers and senior HR executives; (3) investors from U.S.-based venture capital and private equity firms active in HR technology; and (4) academic and thought leaders from U.S.-based research universities and federal professional education institutions. Fifty interviews were completed across 14 industries in North America, Europe, and Asia. Participants were recruited primarily through independent outreach (e.g., professional networks, cold outreach, and academic contacts), supplemented by a limited number of sponsor-facilitated introductions; see the Author Note and Section 4.6 for the full disclosure and discussion of limitations.

4.3. Data Collection

Data were collected through four modalities: virtual meetings conducted via video conferencing platforms (n = 39), in-person interviews (n = 5), messaging-based exchanges (n = 5), and email interviews (n = 1). Virtual and in-person interviews were semi-structured, following the interview protocol in Appendix A while allowing probing and elaboration based on participant responses. Video recordings were produced with participant consent and served as the basis for transcript generation. Participant confidentiality was maintained throughout: organizations are identified by sector and size category rather than by name, and quotations are attributed to role designation rather than to named individuals.

Data collection occurred between August 2025 and January 2026. Analytical preparation, including transcript familiarization, contemporaneous memo consolidation, and finalization of the 15 theory-derived propositions, was conducted concurrently with late-stage data collection. The analytical framework and decision rules were fully specified prior to full-dataset analysis.

Between February and March 2026, AI-assisted retrieval was used to systematically surface proposition-relevant evidence across the complete transcript corpus. Analytical evaluation, classifying evidence as supporting, contradicting, or non-evidentiary, was completed during this period using a predefined decision rule and full-context transcript review. Substantive manuscript drafting commenced only after analytical classification and synthesis were completed. While interpretive exposition was refined during writing, no propositions were added, removed, or reclassified during manuscript preparation. The mean interview duration was 25 minutes (range: 20-35 min). Prior to each interview, participants were informed of the study's purpose and provided verbal consent to audio recording. Interview recordings and transcript files were stored on an institutional server with access restricted to the researcher. In accordance with the University of Notre Dame's Data Retention and Access policy and any applicable sponsor or journal requirements, primary research data (including recordings) will be retained for a reasonable period after publication or project closeout (and, where applicable, at least five years after final fiscal and technical reports) and then securely destroyed.

4.4. Interview Protocol

Two versions of the interview protocol were employed: one for practitioners (HR executives, TA leaders, recruiters) and one for investors, to ensure questions were calibrated to each participant's domain of expertise. The questions were designed to be conversational and hypothesis-probing rather than leading, using open framing to surface authentic practitioner experience rather than confirming pre-existing assumptions. The full protocol is reproduced in Appendix A.

The practitioner protocol centered on six core questions exploring: where AI is currently adding value in TA; how companies are structuring talent acquisition in response to platform-based and integrated approaches; what works and what falls short in operationalizing skills-first hiring; how compliance is shaping executive-level vendor and product decisions; what conversations are occurring internally around regulations such as the EU AI Act and AEDT; and what would make a claim of being an 'AI-first' RPO credible rather than merely marketing.

The investor protocol centered on seven questions exploring: how investment targets navigate the build-vs-buy decision on AI hiring tools; what signals investors pay attention to in a space that is both high-potential and structurally complex; whether recent investments have placed AI at the center of the value thesis; what skills-based hiring looks like in practice in portfolio companies; how founders navigate regulatory compliance as a product design constraint; how investors evaluate efficiency metrics like time-to-hire reduction; and how operators balance flexibility and control in AI-powered TA platforms.

4.5. Data Analysis: AI-Assisted Transcript Analysis

Video recordings were transcribed to generate full interview transcripts. Key insights were noted during and immediately after each interview by the researcher to capture participant intent, emphasis, and context. AI-assisted transcript analysis was then employed to systematically surface themes, patterns, and hypothesis-relevant evidence across the 50-interview dataset, using Airtable's AI functionality (Omni) as an analytical retrieval tool. This approach is increasingly recognized as a legitimate and efficient method for large-sample qualitative research, providing consistency in pattern detection at a scale that manual analysis alone would struggle to maintain across 50 full transcripts.

The AI-assisted retrieval was used to: (a) identify recurring themes and terminology across transcripts; (b) flag hypothesis-relevant passages for each of the 15 propositions tested; and (c) surface representative direct quotations for each theme. The researcher retained full analytical authority over interpretation - AI tools organized and retrieved evidence; the researcher determined meaning, assessed validity, and drew conclusions. For example, the AI repeatedly surfaced passages linking "speed" or "time-to-fill reduction" to caveats about quality controls (e.g., calibration, oversight, and safeguards). The researcher then reviewed the full transcript context for these excerpts, cross-checked them against contemporaneous interview notes, and assessed whether the evidence supported the proposition that AI screening reduces time-to-fill without degrading candidate quality. Only excerpts that remained consistent when read in context, and that aligned with the broader pattern of evidence across interviews, were interpreted as supporting the finding; isolated or ambiguous statements were treated as insufficient for validation.

Each proposition was assessed using a consistent, symmetric decision rule. For each interview, evidence relevant to a proposition was **classified** as **supports**, **contradicts**, or **no clear evidence** after full-context transcript review. Proposition status was then determined by the balance and quality of evidence across the dataset: a proposition was classified as **validated** when supporting evidence was (a) substantively direct (i.e., addressed the claim rather than tangential themes), (b) corroborated across multiple independent interviews and stakeholder types, and (c) more compelling than contradicting evidence when considered in context; it was classified as **invalidated** when the opposite pattern held. Where the dataset contained only limited, ambiguous, or weakly articulated evidence in either direction, the proposition was labeled **emergent**. To support transparency for readers, the paper also reports the number of interviews in which a proposition-relevant supporting instance was identified (n); these counts are presented as descriptive indicators of prevalence and are not treated as statistical confidence thresholds.

Interview tracking and participant management were administered using Airtable, which served as a relational database for outreach sequencing, scheduling, and interview status. Researcher notes and contextual observations recorded during and immediately following each interview were logged directly within the same system, maintaining a unified record linking participant metadata to analytical commentary.

Airtable's AI functionality (Omni) was employed as a retrieval and summarization tool rather than a primary interpreter. The prompting approach was structured and hypothesis-driven: the researcher developed targeted prompts aligned to each of the 15 propositions, directing the tool to surface relevant passages, identify recurring

language patterns, and retrieve candidate quotations. All AI-generated outputs were reviewed, interrogated, and selectively incorporated by the researcher. Where AI retrieval and researcher judgment diverged, researcher interpretation prevailed. This workflow ensured that AI contributed efficiency and consistency at the retrieval stage while analytical authority, including thematic interpretation, validity assessment, and conclusion-drawing, remained with the researcher throughout.

Researcher presence and transcript verification. Because the researcher personally conducted all 50 interviews, either in-person or via video conferencing, primary quality assurance was embedded in the data collection process itself. Contemporaneous notes capturing participant intent, emphasis, and contextual observations were recorded during and immediately following each interview and logged within the same Airtable system used for transcript storage, maintaining a unified record linking participant metadata to analytical commentary. Following each interview, recordings were transcribed to generate full interview transcripts, which were stored in Airtable alongside researcher notes and served as the basis for AI-assisted retrieval. The researcher's direct presence across all interviews, combined with the integration of contemporaneous notes and full transcripts in a single relational database, provided the primary verification mechanism ensuring that AI-retrieved excerpts were evaluated against the researcher's first-hand knowledge of each conversation. All AI-generated outputs were reviewed against this integrated record; where AI retrieval and researcher judgment diverged, researcher interpretation, informed by direct interview experience prevailed.

To support methodological transparency, the AI-assisted workflow was documented in Appendix C, including the database structure used, the query approach aligned to each proposition, and the decision rules applied when evaluating retrieved excerpts. Prompts were structured around each of the 15 propositions and directed Airtable's Omni to surface hypothesis-relevant passages, identify recurring terminology, retrieve candidate quotations, and generate descriptive summaries across the interview dataset. When prompt outputs were reviewed and refined, the researcher recorded the basis for any adjustments to ensure the retrieval strategy remained aligned with the proposition's operational definition. Excerpts were retained as supporting evidence only when they (a) remained substantively consistent when read in the full transcript context, (b) addressed the proposition directly rather than tangentially, and (c) were corroborated by independent instances across interviews; otherwise they were treated as insufficient.

4.6. Limitations

This study carries several limitations. The cross-sectional design precludes causal inference; findings describe patterns at a point in time rather than tracking how adoption outcomes evolve. Self-report data are subject to social desirability bias, particularly among executives discussing regulatory compliance and technology performance. Although the sample includes some recruiters within the practitioner pool (Appendix B), candidates were not sampled, and frontline recruiter perspectives were not analyzed as a dedicated, standalone category. As a result, outcome-oriented claims referenced in the findings (e.g., candidate quality, early attrition, hiring manager satisfaction, and day-to-day recruiter role transformation) should be interpreted as perceptions filtered through organizational stakeholders, often leaders, who may have reputational or professional incentives to report successful adoption. Future research should employ multi-informant designs that triangulate leaders, frontline recruiters, hiring managers, and candidates, and should incorporate direct outcome measures where feasible. Because 3 of 50 participants were recruited via sponsor-facilitated introductions used solely to reach otherwise hard-to-reach practitioner roles, the sample may modestly over-represent perspectives adjacent to the sponsor's professional ecosystem. This risk is mitigated by the small proportion of sponsor-facilitated interviews, the use of consistent inclusion criteria and interview protocols across recruitment channels, and the paper's practice of treating propositions as validated only when corroborated across multiple stakeholder types and independently recruited interviews. The sample, while diverse, underrepresents organizations from the Global South. The use of AI-assisted transcript analysis, while legitimate, introduces the specific tool selection and prompting strategy as methodological variables that are not yet standardized in the literature; to mitigate this limitation and support evaluability, this study documents its AI-assisted workflow, representative prompts, and quality-checking procedures in Appendix C. It should be noted that Guest et al.'s (2006) saturation benchmark of 12 interviews applies to relatively homogeneous samples. Given this study's explicitly heterogeneous design, spanning 14 industries, four stakeholder categories, and three continents, subsequent research suggests that 20 to 40 interviews may be required to achieve metathematic saturation in multi-site, cross-cultural studies (Hagaman & Wutich, 2017). The total of 50 interviews in this study substantially exceeds even this higher threshold, supporting confidence in the breadth of evidence across the dataset.

5. Findings

Findings are organized around four themes that cross-cut the 15 hypotheses tested: (1) AI efficiency gains and their limits; (2) regulatory and compliance barriers; (3) governance and organizational readiness; and (4) market dynamics and the investor perspective. Table 1 presents all 15 hypotheses with validation status and confidence levels. Section 5.5 discusses the three invalidated hypotheses as a distinct analytical contribution.

Table 1 All 15 propositions tested, status (validated/invalidated/emergent), interview support count (n), and confidence level.

Hypothesis	Status (Validated/Invalidated/Emergent*)	n	Confidence
AI in candidate screening reduces time-to-fill	Validated	39	High
Regulated industries concerned about AI bias and compliance	Validated	29	High
TA leaders invest in AI demonstrating time-to-hire reduction without quality loss	Validated	23	High
CHROs hesitant to adopt AI for determinative hiring decisions due to regulations	Validated	22	High
Training recruiters on AI use is a core activity during rollouts	Validated	16	High
Skills-based hiring surfaces more qualified candidates per requisition	Validated	15	High
Executive oversight leads to higher on-time AI implementation rates	Validated	4	Moderate
Investor interest driven by new revenue streams, not cost reduction alone	Validated	3	Moderate
Senior leader expectations increase recruiter AI adoption rates	Validated	3	Moderate
HR teams linking AI to retention receive faster project approval	Emergent	2	Lower
AI talent initiatives in OKRs reach measurable impact faster	Emergent	1	Lower
Reputational risks of flawed AI implementation outweigh first-mover benefits	Emergent	1	Lower
Clients willing to pay premium for inclusive AI outcomes	Invalidated	3	Moderate
Clients willing to pay premium for AI speed outcomes	Invalidated	5	Moderate
Clients adopt tools at higher rates if unbundled from RPO contracts	Invalidated	6	Moderate

Note: In this table, **n** denotes the number of interviews containing **supporting** evidence for the proposition. The *n* values are reported as descriptive indicators of how widely a pattern was observed in this sample; they are not treated as statistical confidence thresholds. Status reflects the net balance and qualitative strength of evidence across interviews: propositions were labeled **validated** when supporting evidence was more compelling and more prevalent than contradicting evidence, **invalidated** when contradicting evidence was more compelling and more prevalent than supporting evidence, and **emergent** when evidence in either direction was limited or ambiguous. *Emergent propositions are presented as hypotheses for future empirical investigation rather than as validated findings.

5.1. AI Efficiency Gains and Their Limits

The most robustly validated finding confirmed across 39 of 50 interviews is that AI-assisted candidate screening reduces time-to-fill rates. This held consistently across company sizes, industries, and geographies, suggesting it represents a near-universal operational reality for organizations that have implemented AI screening at meaningful scale. The WEF Future of Jobs Report 2025 contextualizes why this finding matters competitively: as 85% of employers plan upskilling as their primary workforce strategy and skill gaps affect 63% of organizations, the ability to identify qualified candidates at speed is a decisive operational advantage (WEF, 2025).

The primary mechanism identified across interviews is automation of high-volume, low-judgment tasks - resume parsing, initial ranking, scheduling, and administrative communication, freeing recruiters to focus on relationship-building and final evaluation. However, participants consistently qualified this narrative: speed gains were described as dependent on the quality of criteria encoded into AI systems and on the presence of quality safeguards.

Several participants also emphasized that poorly governed implementations can generate negative downstream consequences, including elevated early attrition and hiring manager dissatisfaction. These accounts reflect participant perceptions based on organizational experience rather than direct measurement from candidates or hiring managers within this study. These accounts were most commonly associated with AI deployments optimized primarily for speed rather than for quality, calibration, and oversight.

A related high-confidence finding (23 interviews) confirms that TA leaders specifically require evidence that AI reduces time-to-hire without negatively impacting candidate quality before committing to investment. This dual requirement: faster and better, distinguishes successful implementations from those that generate efficiency metrics without sustained organizational value.

"AI significantly reduces the time-to-fill rate. The manager estimated it saved him about 26 hours of manual time in a single month by automating tasks."

— **Talent Acquisition Manager, Information Technology & Services**

"AI handles first-round screening, typically narrowing 100 applicants to top 10. Moving toward AI-based initial candidate interviews before human involvement."

— **Senior Talent Acquisition Leader, Staffing & Recruiting (Civil Engineering focus)**

"AI is doing the heavy lifting on the first pass — our recruiters now spend their time on the conversations that actually matter."

— **VP of Talent Acquisition, Enterprise Technology Firm**

"AI is used to prioritize individuals who most match the minimum requirements, not to disqualify those who are close."

— **Director of Talent Acquisition, Information Technology & Services**

5.2. Regulatory and Compliance Barriers

Across 29 interviews with particular concentration in financial services, healthcare, and government, AI bias and regulatory compliance emerged as the dominant barriers to adoption beyond administrative efficiency use cases. The EU AI Act and NYC's AEDT were the most frequently cited instruments, reflecting the regulatory landscape's rapid formalization. A related high-confidence finding (22 interviews) documents that CHROs and senior HR leaders express active hesitancy specifically toward AI in determinative hiring decisions, distinct from hesitancy toward AI broadly. Most participants expressing this hesitancy simultaneously reported enthusiasm for AI in administrative TA functions.

Across interviews, participants most frequently cited OFCCP (Office of Federal Contract Compliance Programs), GDPR (General Data Protection Regulation), HIPAA (Health Insurance Portability and Accountability Act), SOC 2 (Service Organization Control), AEDT/ADS (Automated Employment Decision Tools - NYC & California), and EEO (Equal Employment Opportunity) requirements as the key compliance frameworks shaping AI-enabled hiring decisions. The most common barrier categories were data privacy, bias/fairness risk, general litigation exposure, sector-specific regulatory constraints, and auditability requirements.

"There is far too much pending litigation on the use of AI in candidate review to consider it. We will not even partner with agencies who use it in that regard." - Senior HR Leader, Regulated Financial Services Organization

This pattern reflects what this study terms *anticipatory compliance behavior*: organizations slowing, constraining, or deferring AI-enabled hiring use cases in response to regulatory signals and perceived future enforcement risk, even in advance of formal audits or penalties. The EU AI Act's extraterritorial scope (European Commission, 2024) and the unresolved tension between the Act's bias-detection provisions and GDPR's data processing restrictions (De Luca & Federico, 2025) together produce regulatory ambiguity that practitioners cannot easily resolve through conventional legal review processes.

Analytically, anticipatory compliance behavior operates as an adoption "brake" driven less by direct legal prohibition than by uncertainty about evidentiary standards (e.g., what constitutes an acceptable bias audit, documentation threshold, or lawful basis for processing sensitive attributes for bias detection). As a result, organizations tend to narrow AI use to low-risk administrative tasks, delay expansion into determinative decision points, and elevate risk, legal, and security stakeholders as de facto gatekeepers of product and vendor selection. For policymakers, this suggests that clearer guidance and interoperable standards may accelerate beneficial adoption without weakening protections.

Table 2 Tally of interviews on regulatory and compliance barriers by industry.

Industry Sector	Number of Interviews	Key Compliance Risk Driver
Information Technology & Services	6	Sector-Specific (Health/Finance and Government)
Financial Services	4	GDPR
Staffing & Recruiting	4	General Litigation
Hospital & Health Care	2	Data Privacy
Venture Capital	2	Bias/Fairness
Human Capital Services	2	Data Privacy
Medical Devices	1	Sector-Specific (Health/Finance)
Business Software	1	Auditability
Insurance	1	Data Privacy
Human Resources	1	Bias/Fairness
Government Relations	1	General Litigation
Higher Education	1	Disclosure/Transparency
Retail	1	Other
Construction	1	Other
Consulting	1	Data Privacy

Note: For cross-sector service providers (e.g., Information Technology & Services), the listed risk driver reflects the primary regulated client context or use case referenced in interviews (e.g., health, finance, government), rather than the provider's sector classification.

"The biggest hindrance to using AI during critical stages like interviews is security litigation and regulatory compliance. The company works with government clients and must comply with strict standards like SOC 2 and HIPAA"

— Senior HR Leader, Information Technology & Services (Government Sector)

5.3. Governance Structures and Organizational Readiness

Four validated hypotheses converge on a consistent finding: AI initiatives with formal executive sponsorship, C-suite visibility, and OKR alignment consistently outperform grassroots adoption efforts on implementation timeline, recruiter adoption rate, and measurable business impact. Institutional Theory explains this pattern: formal organizational endorsement institutionalizes AI adoption as an expected practice rather than a voluntary individual choice, transforming the adoption environment in ways that individual-level TAM constructs alone cannot predict.

Recruiter training emerged as an equally high-confidence finding (16 interviews), with participants consistently describing training as constitutive of, not supplementary to, AI rollouts. Effective programs were described as covering

three domains: technical fluency (how the tool works), ethical literacy (how bias enters AI systems and how to detect it), and judgment calibration (when to trust AI output and when to override it). Participants who described high adoption rates consistently cited training investment as a prerequisite.

"Training is essential to ensure individuals do not solely rely on AI. Training should also cover the ethical use of AI, the cons and biases it presents, and maintaining human integrity."

— HR Leader, Interview Participant

5.3.1. Governance structures and executive sponsorship mechanisms

Participants described governance as a practical mechanism for de-risking AI adoption rather than a purely administrative overlay. In the subset of interviews where training and governance were discussed explicitly (n = 16), participants consistently emphasized that adoption was most durable when AI initiatives were framed as enterprise priorities with clear decision rights, rather than as discretionary experimentation within TA. Where executive oversight was present, participants described faster movement from pilot to implementation and fewer internal veto points, aligning with the moderate-confidence hypothesis that executive oversight supports on-time implementation (n = 4) and that senior leader expectations increase recruiter adoption (n = 3).

A common pattern was top-down executive sponsorship, particularly in large organizations, in which C-suite cost, productivity, or standardization goals established an explicit mandate for AI exploration and rollout. In these cases, participants reported that individual functions were expected to translate the mandate into goal-oriented business cases and implementation plans that were traceable to enterprise objectives (e.g., efficiency metrics, service-level targets, or operational OKRs). This framing connects to the validated finding that AI initiatives embedded in OKRs can reach measurable impact faster (n = 1), suggesting that formal performance systems can function as an enabling infrastructure for adoption.

In highly regulated settings, participants described hybrid approval processes combining top-down alignment (often involving CIO/technology leadership, legal, and risk stakeholders) with bottom-up pilot activity led by TA teams. This structure was presented as a way to preserve strategic oversight and compliance while still enabling local experimentation to establish fit-for-purpose workflows. Relatedly, some participants referenced standing AI councils or committees, typically cross-functional groups spanning HR, legal, and technology, responsible for reviewing use cases, establishing guidance, and in some cases coordinating training and change management.

Finally, participants in risk-averse contexts described policy-driven governance in which formal guidelines clarified where AI could be used (e.g., administrative workflow support vs. determinative decision-making) and what documentation or review steps were required. These governance approaches were described as complementary to external compliance demands because they operationalized internal accountability, including expectations for transparency, audit trails, and appropriate human oversight.

5.3.2. Bottom-up failure modes in the absence of executive sponsorship

In contrast, participants described bottom-up adoption efforts as vulnerable to stall-out when they lacked executive sponsorship, particularly when cross-functional dependencies (e.g., IT security review, procurement, or legal sign-off) were required. Several participants characterized bottom-up proposals as more likely to be deprioritized due to resource constraints or misalignment with enterprise goals, even when local TA teams perceived clear operational value. Where centralized governance was absent, participants also reported "tool sprawl," in which teams adopted multiple uncoordinated AI tools that did not integrate with each other or with core TA systems, producing inefficiencies and fragmented workflows.

Participants further described the compliance and security risks that can accompany informal adoption. One incident recounted involved use of an unauthorized AI tool to record proprietary meeting content, with data stored in an uncontrolled cloud environment; the event was described as a governance "wake-up call" that accelerated the shift toward sanctioned, mainstream tools designed to control data flow and reduce exposure. In these accounts, governance was not positioned as an inhibitor of innovation but as the condition that made scaled adoption acceptable to risk owners.

5.3.3. Measurable outcomes attributed to training investment

Consistent with the high-confidence training finding (n = 16), participants linked training investment to both adoption and operational outcomes. Reported impacts included time savings and efficiency gains (e.g., reduced manual work

hours following automation), improved candidate communication and engagement through repeated enablement sessions (e.g., weekly internal “crash courses” on tool use), and upskilling that supported internal mobility as practitioners acquired AI-relevant capabilities. Participants also emphasized the importance of role-specific training, noting that tailoring enablement to distinct recruiting functions (e.g., sourcing vs. interviewing) increased effective usage; for example, by supporting the generation of structured interview questions or enabling more targeted candidate search strategies.

5.4. Market Dynamics and the Investor Perspective

Interviews with venture capitalists and investors revealed a perspective on AI in TA that is largely absent from the existing academic literature. Three validated hypotheses, albeit at moderate to lower confidence levels, concern investor orientation and map to the competitive-level dynamics captured by the Resource-Based View (RBV): investors prioritize new, monetizable revenue stream creation over cost reduction when evaluating AI-powered TA companies; they treat reputational risks of flawed AI implementation as significant barriers to investment; and they require evidence of compliance readiness as a baseline condition.

Investor interview questions specifically probed the build-vs-buy decision, signals of genuine vs. branded AI capability, and how efficiency metrics like time-to-hire reduction are validated. The consistent signal from investor participants was skepticism toward efficiency claims without documented validation methodology, and preference for companies demonstrating defensible differentiation - proprietary data, workflow integration, or demonstrated compliance architecture - over those making broad capability claims.

This investor emphasis on “defensible differentiation” aligns directly with the Resource-Based View (RBV) framing introduced in Section 3. From an RBV perspective, investors were effectively evaluating whether an AI-enabled TA capability meets Barney’s (1991) criteria for sustained advantage: whether it is *valuable* (e.g., produces measurable efficiency or decision-quality gains), *rare* (e.g., access to proprietary or uniquely permissioned data), and *difficult to imitate* (e.g., embedded workflow integration, defensible models, and compliance architectures that are costly to replicate). In this sense, investor skepticism toward generic “AI claims” can be interpreted as skepticism toward non-VRIN resources that competitors can rapidly copy without sustained advantage.

When asked to specify evaluation criteria for AI-powered TA companies, investor participants emphasized that diligence is anchored less in aspirational product narratives than in demonstrated, externally verifiable outcomes. In particular, investors reported prioritizing evidence of validated efficiency gains (e.g., time-to-hire reduction or cost savings) supported by customer references, with validation conducted through direct user conversations and reference checks rather than relying on vendor-reported metrics alone. Participants noted that differentiation is a critical screening lens given market saturation, and they described genuine AI capability as distinguishable from marketing when companies can substantiate defensible advantage through demonstrable product performance, sustained organizational value beyond one-off efficiency claims, and unique features that are difficult to replicate. Finally, investors framed compliance readiness as a baseline condition for investability in this category, including the presence of formal policies, operational safeguards, and, in some cases, risk-transfer mechanisms (e.g., insurance) to manage regulatory exposure and implementation risk.

5.5. Invalidated Hypotheses: What the Data Refutes

Three propositions were not supported by the interview evidence. Each carries direct implications for practitioners and vendors. Consistent with the decision rule in Section 4.5, these propositions are labeled invalidated because contradicting evidence was more prevalent than supporting evidence across interviews, even though some attracted a moderate number of supporting endorsements.

First, the hypothesis that clients would pay a premium for AI features improving inclusive outcomes was invalidated across three interviews. Participants consistently treated inclusive outcomes as a baseline expectation - a regulatory and reputational floor, rather than a differentiating feature warranting premium pricing. This does not imply that inclusive outcomes are unimportant; it reflects that the market has not yet developed the auditing and measurement infrastructure to verify and communicate inclusive outcomes with sufficient credibility to support price differentiation.

Second, the hypothesis that clients would pay a premium for AI features improving speed outcomes was invalidated across five interviews. Speed gains are increasingly treated as table stakes rather than differentiators. This aligns with Everest Group’s (2024) observation that market differentiation among RPO providers has shifted from efficiency metrics - now expected as hygiene - toward consulting, advisory, and skills-based capabilities.

Third, the hypothesis that clients would adopt workflow-enablement tools at higher rates if unbundled from full-service RPO contracts was invalidated across six interviews. Large organizations consistently preferred integrated, enterprise-level solutions over point solutions - directly contradicting a common vendor narrative. This finding is consistent with the Everest Group's (2025) PEAK Matrix characterization of market leaders as comprehensive, integrated platform providers.

6. Discussion

6.1. Theoretical Contributions

This study makes three theoretical contributions. First, it provides cross-industry, cross-geographic empirical support for the integrated TAM-Institutional Theory framework, confirming that individual-level adoption attitudes are necessary but insufficient, organizational-level institutional factors frequently determine outcomes. The governance finding is particularly confirmatory: C-suite-sponsored AI initiatives outperform bottom-up efforts because institutional endorsement transforms perceived usefulness from an individual calculation into an organizational expectation.

Second, by situating organizational findings within Everest Group market data and WEF Future of Jobs evidence, the study establishes coherence between practitioner-level interview observations and macro-level structural dynamics. This cross-validation increases confidence that the patterns identified are structural rather than idiosyncratic.

6.2. Methodological Contribution: AI-Assisted Transcript Analysis

This study contributes a methodological observation relevant to large-sample qualitative research: when the researcher retains primary interpretive authority and is personally present across all data collection, AI-assisted database querying can enhance retrieval efficiency and consistency without displacing analytical judgment. In this study, all 50 interviews were conducted directly by the researcher, generating contemporaneous notes and full transcripts stored within a structured relational database. Airtable's AI functionality (Omni) was then used to query that database, surfacing hypothesis-relevant passages, identifying recurring terminology, retrieving candidate quotations, and generating descriptive summaries across the dataset. Critically, the researcher's first-hand knowledge of each interview provided the evaluative frame within which all AI-retrieved outputs were assessed; passages were retained only when they were consistent with the researcher's direct recall and contemporaneous notes, addressed the proposition substantively, and were corroborated across multiple interviews. This approach is itself consistent with the augmentation model the research validates: AI expands retrieval capacity and supports consistency at scale; human judgment, grounded here in direct interview experience, determines meaning.

6.3. Practical Implications

For HR and TA practitioners, four implications flow directly from the findings. First, prioritize low-risk use cases where value can be demonstrated quickly (e.g., workflow coordination and administrative automation) before expanding into determinative decision points. Second, treat compliance and documentation requirements as design constraints from day one rather than as post-deployment add-ons. Third, invest in structured enablement so recruiters understand both tool operation and appropriate limits of reliance. Fourth, secure executive sponsorship and clear decision rights before scaling, particularly where cross-functional approvals (legal, security, procurement) create potential veto points.

For RPO providers and HR technology vendors, the invalidation findings are commercially significant. Speed-related features are increasingly treated as table stakes rather than differentiators, suggesting limited pricing power for "faster hiring" claims absent externally verifiable value. Product strategy should also reflect enterprise preference for integrated solutions rather than unbundled point tools. More broadly, vendors should assume that buyers will expect inclusion and compliance capabilities as baseline requirements and will scrutinize whether those claims are auditable and defensible.

For policymakers, the prevalence of *anticipatory compliance behavior* suggests that clarity is as important as stringency: regulatory ambiguity can slow even beneficial, governance-aware adoption. The unresolved AI Act-GDPR tension (De Luca & Federico, 2025) is producing conservatism that may constrain beneficial applications alongside harmful ones.

6.4. The Human-AI Partnership in Talent Acquisition

Across the 50 interviews, participants overwhelmingly framed AI as augmenting rather than replacing human judgment in talent acquisition. This broad direction is consistent with macro-level accounts such as the World Economic Forum's (2025) emphasis on human-machine collaboration; however, the interview evidence adds practitioner-level nuance

that is largely absent from employer-survey summaries. In particular, the data refine the augmentation narrative by identifying five recurring “speed bumps” and non-negotiables: a hype-versus-reality gap in what is marketed as AI, regulatory fear as an active adoption brake in determinative decisions, economic resistance to premium pricing for “inclusive AI” and speed features, a shift in recruiter identity toward human integrity and judgment rather than capacity alone, and the decisive role of top-down versus bottom-up adoption dynamics within enterprises.

First, the interviews surfaced a clear “hype vs. reality” gap in tooling. Participants repeatedly distinguished between basic automation labeled as “AI” and genuinely AI-enabled functionality applied to complex decision-making. In practice, augmentation was described as most mature in administrative and coordination tasks (e.g., parsing, scheduling, templated communications, and note-taking), aligning with the efficiency mechanism described in Section 5.1. By contrast, machine-led selection at determinative stages was described as comparatively rare and, in many contexts, strategically avoided.

Second, regulatory fear functioned as an active brake, not merely a background barrier. Beyond general concerns about “regulation,” senior leaders described specific hesitation tied to the EU AI Act and AEDT requirements when AI use cases approached determinative decision points (Section 5.2). Notably, the study documented *anticipatory compliance behavior* - organizations slowing, constraining, or deferring AI-enabled hiring use cases in response to perceived enforcement risk, even prior to audits or penalties. In one account, a regulated-industry leader stated they would not partner with agencies using AI for candidate review due to pending litigation risk, illustrating how liability concerns can actively stall augmentation trajectories that survey-based accounts are unlikely to capture.

Third, the interviews complicate optimistic “augmentation-for-DEI” narratives by highlighting an economic constraint: the market is not yet willing to pay a premium for speed or inclusion features. Section 5.5 shows that hypotheses about premium pricing for inclusive AI outcomes and for speed improvements were invalidated. Practitioners described inclusive outcomes as a baseline expectation, often treated as a regulatory and reputational floor, while simultaneously resisting additional fees for the measurement, auditing, and model-development costs required to substantiate inclusion claims. The implication is that, even where augmentation could support more inclusive hiring, adoption and product investment may lag unless buyers are willing to fund the infrastructure that makes inclusive outcomes credible and verifiable.

Fourth, participants framed augmentation as a redefinition of professional identity: the recruiter’s value is increasingly tied to maintaining human integrity and exercising judgment in areas where AI is perceived as weak (e.g., culture fit, soft skills, and critical thinking), rather than simply “doing more with less.” This interpretation gives analytical weight to the training finding in Section 5.3, where effective programs were described as emphasizing ethical literacy and judgment calibration to prevent overreliance on tools. It also helps explain why poorly governed deployments can produce negative downstream consequences (Section 5.1): when augmentation is interpreted as permission to defer judgment, quality can degrade even as efficiency metrics improve.

Fifth, the interviews underscored that augmentation is mediated by internal power dynamics as much as by technical capability. As shown in Section 5.3, top-down sponsorship and OKR alignment reduce veto points and accelerate movement from pilot to scaled implementation, whereas bottom-up adoption efforts must repeatedly justify economic value and are more vulnerable to stall-out amid security, procurement, and legal reviews. Accordingly, “human-AI partnership” in TA is better understood as an organizational negotiation over cost, compliance, and accountability than as a uniform workforce strategy that spreads evenly across functions.

The skills-based hiring finding provides a concrete illustration of augmentation under these constraints: AI does not make human judgment redundant in candidate evaluation; it changes where judgment is applied by surfacing a broader pool for human evaluators to assess and by structuring evidence about skills more consistently. This reconfiguration of the human-AI division of labor in TA has implications for recruiter role design, training investment, and professional identity, particularly as organizations navigate the tradeoffs among efficiency, governance, and compliance risk.

7. Conclusion

This study examined the organizational, regulatory, and governance dimensions of AI adoption in global talent acquisition through a qualitative, hypothesis-driven interview design comprising 50 structured interviews spanning 14 industries in North America, Europe, and Asia, with data captured via video interviews, recordings and transcripts analyzed using AI-assisted tools. Its core contributions are threefold: it establishes with high-confidence evidence that AI integration reduces time-to-fill but that this gain is contingent on quality safeguards and governance structures many organizations lack; it situates organizational findings within market evidence from Everest Group and WEF Future of

Jobs 2025, establishing coherence between practitioner experience and structural market dynamics; and it demonstrates through systematic hypothesis invalidation that premium pricing for speed and inclusion features is not currently supported, and that integrated solutions are preferred over unbundled tools, findings with direct strategic implications for vendors and investors.

Future research should address the cross-sectional limitation through longitudinal designs tracking how adoption outcomes evolve from pilot to scale; expand geographic sampling to include the Global South; develop multi-informant designs that triangulate leader, recruiter, and candidate perspectives; and establish reporting standards for AI-assisted transcript analysis as a qualitative research method. As the WEF projects 170 million new jobs created globally by 2030 alongside 92 million displaced, the talent acquisition function that connects people to those opportunities carries stakes that justify sustained, rigorous scholarly attention.

Compliance with ethical standards

Acknowledgments

The author gratefully acknowledges the 50 industry practitioners, executives, and investors who generously gave their time to participate in this study, and Cielo for providing the industry context that made this research possible.

Disclosure of conflict of interest

This research was conducted as part of a graduate capstone project at the University of Notre Dame's ESTEEM Program with industry support. The sponsor's support was limited to providing contextual industry access and licensed market research resources; the sponsor did not direct the research questions, control data collection, influence analysis, or review findings prior to submission. The author declares no financial conflict of interest. For additional transparency regarding sponsor involvement and sampling, see the Author Note and Section 4.6.

Statement of Ethical Approval

This research was conducted in accordance with the ethical principles of the University of Notre Dame and the guidelines set forth in the Belmont Report (1979). The study involved semi-structured interviews with adult professionals on non-sensitive workplace and business topics, with no collection of personally identifiable sensitive information and no foreseeable risk of harm to participants. The author determined that this research meets the criteria for exemption under 45 CFR 46.104(d), Exemption Category 2.

Statement of informed consent

Informed consent was obtained from all individual participants included in the study. Prior to each interview, participants were informed of the voluntary nature of their participation, the purpose of the research, how their responses would be used, the anonymization procedures applied to all data, and their right to withdraw at any time without consequence. For virtual and in-person interviews, verbal informed consent was obtained and, where applicable, recording consent was secured before any recording commenced. No participant is identified by name in this manuscript; all quotations are attributed to role designation and industry category only.

Funding

This research was supported by an industry sponsor. The sponsor's role was limited to facilitating access to industry context and premium licensed data. The sponsor did not define the research questions, influence the interview protocol, participate in data collection or analysis, review transcripts, review findings, or have any editorial control over the manuscript prior to submission. All analytical decisions, interpretation of findings, and conclusions were made independently by the author. For transparency regarding data provenance and sampling, 3 of 50 interview participants were recruited through sponsor-facilitated introductions used solely to obtain access to otherwise hard-to-reach practitioner roles; all other participants were recruited through independent outreach. The sponsor was not present for interviews and did not have access to recordings, transcripts, participant-level responses, or interim findings.

Data Availability Statement

The Everest Group reports cited in this paper (EGR-2024-26-R-6842 and EGR-2025-26-R-7224) are proprietary commercial research publications accessed through a licensed industry research engagement conducted as part of the ESTEEM capstone program at the University of Notre Dame. Because these reports are subject to third-party licensing

restrictions, they cannot be publicly deposited with this manuscript and may not be shareable in full text with reviewers or readers who are not licensed subscribers. Accordingly, Everest Group materials are used as contextual market intelligence (e.g., to describe broad RPO market conditions and positioning trends) rather than as independent corroborating evidence for the study's interview-based findings. Readers with Everest Group access may consult the reports directly using the identifiers above; any sharing beyond citation metadata is subject to Everest Group's licensing terms. The author will align the final data availability and transparency disclosures with the target journal's policy at the time of submission.

References

- [1] Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99-120. <https://doi.org/10.1177/014920639101700108>
- [2] Becker, G. S. (1964). Human capital: A theoretical and empirical analysis, with special reference to education. *National Bureau of Economic Research*.
- [3] Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319-340. <https://doi.org/10.2307/249008>
- [4] De Luca, S., & Federico, M. (2025). Algorithmic discrimination under the AI Act and the GDPR (PE 769.509). *European Parliamentary Research Service*. [https://www.europarl.europa.eu/RegData/etudes/ATAG/2025/769509/EPRS_ATA\(2025\)769509_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/ATAG/2025/769509/EPRS_ATA(2025)769509_EN.pdf)
- [5] DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147-160. <https://doi.org/10.2307/2095101>
- [6] European Commission. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- [7] Everest Group. (2024). *Recruitment Process Outsourcing (RPO) Services - State of the Market 2024* (Report No. EGR-2024-26-R-6842). Everest Group Research. Available to licensed subscribers at <https://www.everestgrp.com>
- [8] Everest Group. (2025). *Recruitment Process Outsourcing (RPO) Services PEAK Matrix® Assessment 2025 - Global* (Report No. EGR-2025-26-R-7224). Everest Group Research. Available to licensed subscribers at <https://www.everestgrp.com>
- [9] Ezenwa, C. L. (2025). Artificial Intelligence for Inclusive Innovation: Investigating how AI-driven Information Systems can democratize access to opportunities in underserved communities. *World Journal of Advanced Research and Reviews*, 27(03), 897-908.
- [10] Ezenwa, C. L. (2025). From Hubs to Economies: Rethinking Innovation Ecosystem Models for U.S.-Africa Collaboration. *International Journal of Science and Research Archive*, 17(01), 212-224.
- [11] Gilgun, J. F. (2019). Deductive qualitative analysis and grounded theory: Sensitizing concepts and hypothesis-testing. In A. Bryant & K. Charmaz (Eds.), *The SAGE handbook of current developments in grounded theory* (pp. 107-122). SAGE. <https://doi.org/10.4135/9781526485656>
- [12] Guest, G., Bunce, A., & Johnson, L. (2006). How many interviews are enough? An experiment with data saturation and variability. *Field Methods*, 18(1), 59-82. <https://doi.org/10.1177/1525822X05279903>
- [13] Herold, M., & Roedenbeck, M. R. H. (2025). AI-driven research in the recruitment and selection process: Application of an AI taxonomy with a systematic literature review. *SAGE Open*, 15(3), 1-21. <https://doi.org/10.1177/21582440251361746>
- [14] Hagaman, A. K., & Wutich, A. (2017). How many interviews are enough to identify metathemes in multisited and cross-cultural research? Another perspective on Guest, Bunce, and Johnson's (2006) landmark study. *Field Methods*, 29(1), 23-41. <https://doi.org/10.1177/1525822X16640447>
- [15] Hmoud, B., & Várallyai, L. (2019). Will artificial intelligence take over human resources recruitment and selection? *Network Intelligence Studies*, 7(13), 21-30.

- [16] Mori, M., Sasseti, S., Cavaliere, V., & Bonti, M. (2025). A systematic literature review on artificial intelligence in recruiting and selection: A matter of ethics. *Personnel Review*, 54(3), 854-878. <https://doi.org/10.1108/PR-03-2023-0257>
- [17] National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. (1979). *The Belmont report: Ethical principles and guidelines for the protection of human subjects of research*. U.S. Department of Health, Education, and Welfare. <https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report/index.html>
- [18] New York City Local Law 144 of 2021. Automated Employment Decision Tools Act. New York City Council. <https://www.nyc.gov/site/dca/about/automated-employment-decision-tools.page>
- [19] Paramita, D., Okwir, S., & Nuur, C. (2024). Artificial intelligence in talent acquisition: Exploring organisational and operational dimensions. *International Journal of Organizational Analysis*, 32(11), 108-131.
- [20] Patton, M. Q. (2002). *Qualitative research and evaluation methods* (3rd ed.). SAGE Publications.
- [21] Roppelt, J. S., Schuster, A., Greimel, N. S., Kanbach, D. K., & Sen, K. (2025). Towards effective adoption of artificial intelligence in talent acquisition: A mixed method study. *International Journal of Information Management*, 82, Article 102870. <https://doi.org/10.1016/j.ijinfomgt.2025.102870>
- [22] 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. <https://doi.org/10.1177/0008125619867910>
- [23] World Economic Forum. (2025). The Future of Jobs Report 2025. *World Economic Forum*. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>

8. Appendix A: Interview Protocol

Two protocol versions were used: one for practitioners and one for investors. Questions were designed to be conversational and open-ended, probing hypothesis-relevant themes without leading participants. All virtual and in-person interviews were recorded with participant consent and transcribed for analysis.

Practitioner Protocol (HR Executives, TA Leaders, Recruiters)

The following six questions formed the core of practitioner interviews:

- I've noticed a lot of people talk about AI and hiring, but it's not always clear where it's actually adding value. Where do you see it making the biggest impact right now?
- Some say platform-based hiring is the future; others say it fragments the process. What kind of shifts are you actually seeing in how companies structure talent acquisition?
- With skills-first hiring becoming more popular, I'm curious what you've seen work well or fall short when companies try to operationalize it.
- It feels like compliance is starting to influence product and vendor choices a lot more than it used to. How do you see that shaping decisions at the exec level?
- When it comes to regulations like the EU AI Act or AEDT, what kind of conversations are happening behind the scenes that most people aren't seeing?
- I've seen a few new players claim they're 'AI-first' RPOs. What would make you believe that's more than just branding?

Investor Protocol (VCs and Private Equity)

The following seven questions formed the core of investor interviews:

- I've seen some teams struggle to decide whether to build or buy AI-based hiring tools. How do people you've worked with typically handle that?
- Some founders I've talked to say investors are split on whether the RPO space is ripe for disruption or too messy. What kind of signals do you pay attention to in spaces like this?
- Have you made any recent investment in this sector that leverages AI? If yes, was the AI center a critical consideration?

- A lot of people seem excited about skills-based hiring, but I'm trying to understand what that actually looks like in practice. Have you seen any portfolio companies doing something interesting here?
- It seems like regulatory compliance is starting to drive product decisions more than features. How do the founders you work with navigate that?
- We've seen some companies claim their AI cuts hiring time by 30%, but I'm curious how investors think about validating those kinds of metrics.
- When a product promises both flexibility and control, there's often a tradeoff. I'm trying to understand which side matters more to operators; have you seen any patterns?

9. Appendix B: Author Data Input - Analyzed Interview Evidence

Participant Profile Table (Anonymized)

Table 3 Participants Profile Table showing Title, Industry and Interview style.

Role/Title Category	Industry Sector	Organization Size	Interview Modality	Recruitment Channel
President	Consulting	—	Virtual Meeting	Independent recruitment
Senior Recruiter	Information Services	—	Virtual Meeting	Independent recruitment
Director, Employer Engagement	Higher Education	10,001+ Employees	In-Person	Independent recruitment
Talent Partner	Hospital & Health Care	—	Messaging	Independent recruitment
CEO	Investment	51-200 Employees	Virtual Meeting	Independent recruitment
Talent & Client Engagement Director	Consulting	1001-5000 Employees	Virtual Meeting	Independent recruitment
Director of Government Sales	Government Relations	—	Virtual Meeting	Independent recruitment
Senior Talent Acquisition Advisor	Business Software	501-1000 Employees	Virtual Meeting	Sponsor-facilitated
Director, Talent Acquisition	Automobile	—	Virtual Meeting	Independent recruitment
Culture and DEI Manager	Financial Services	1001-5000 Employees	Virtual Meeting	Independent recruitment
Recruiter	Staffing & Recruiting	—	Virtual Meeting	Independent recruitment
Director of Strategic Insights & Innovation	Human Capital Services	11-50 Employees	Virtual Meeting	Independent recruitment
Partner	Investment	11-50 Employees	In-Person	Independent recruitment
Senior Recruiter, Creative & Marketing	Staffing & Recruiting	—	Messaging	Independent recruitment
Founder	Human Resources	—	Messaging	Independent recruitment

Principal Consultant	Staffing & Recruiting	11-50 Employees	Virtual Meeting	Independent recruitment
Director, Talent Acquisition Events, Marketing, & Inclusion	Financial Services	—	Virtual Meeting	Independent recruitment
Manager, Talent Acquisition	Marketing & Advertising	—	Virtual Meeting	Independent recruitment
Talent Acquisition Manager	Consulting	11-50 Employees	Virtual Meeting	Independent recruitment
Senior Talent Acquisition Specialist	Information Technology & Services	—	Virtual Meeting	Independent recruitment
Sr. Director, Global Talent Acquisition	Technology company	—	Virtual Meeting	Sponsor-facilitated
Senior Executive / IT Recruiter	Staffing & Recruiting	—	Virtual Meeting	Independent recruitment
Corporate Recruiter Lead, Banking Officer	Financial Services	—	Email	Independent recruitment
Senior Recruiter	Construction	501-1000 Employees	Messaging	Independent recruitment
Recruiter / Sourcer	Consulting	11-50 Employees	Virtual Meeting	Independent recruitment
Director Talent Acquisition	Civil Engineering	—	Virtual Meeting	Independent recruitment
Talent Acquisition Manager	Construction	—	Virtual Meeting	Independent recruitment
Recruiter	Hospital & Health Care	—	Virtual Meeting	Independent recruitment
Talent Recruiter	—	—	Virtual Meeting	Independent recruitment
Talent Acquisition	Higher Education	—	Virtual Meeting	Independent recruitment
Senior Recruiter	Information Technology & Services	—	Virtual Meeting	Independent recruitment
Founder/CEO	Venture Capital & Private Equity	—	Virtual Meeting	Independent recruitment
Strategic Sourcing Recruiter	Hospital & Health Care	1001-5000 Employees	Virtual Meeting	Independent recruitment
Talent Acquisition Manager	Medical Devices	—	Virtual Meeting	Independent recruitment
Senior Talent Acquisition Specialist	Information Technology & Services	—	Virtual Meeting	Independent recruitment
Senior Manager, Cloud Engineering	Information Technology & Services	10,001+ Employees	In-Person	Independent recruitment
Recruiter	Financial Services	—	Virtual Meeting	Independent recruitment
Senior HR Specialist	—	—	Virtual Meeting	Independent recruitment

Founder	Software Development	51-200 Employees	In Person	Independent recruitment
Corporate Human Resources	Retail	—	Messaging	Independent recruitment
Human Resources Manager	Manufacturing	—	Virtual Meeting	Independent recruitment
Senior Vice President, Chief Talent Officer	Hospital & Health Care	—	Virtual Meeting	Sponsor-facilitated
Talent Partner	Information Technology & Services	—	Virtual Meeting	Independent recruitment
Founder & CEO	Staffing & Recruiting	11-50 Employees	Virtual Meeting	Independent recruitment
TA Program Manager	Hospital & Health Care	10,001+ Employees	Virtual Meeting	Independent recruitment
Partner, Head of Platform	Venture Capital & Private Equity	11-50 Employees	Virtual Meeting	Independent recruitment
Professor	Higher Education	201-500 Employees	In Person	Independent recruitment
Talent Acquisition Specialist	Information Technology & Services	—	Virtual Meeting	Independent recruitment

Note: Em dashes (—) indicate that organization size and/or industry sector data were not collected or were withheld for confidentiality where disclosure could reasonably identify participants. Recruitment channel indicates whether the interviewee was recruited via independent recruitment (e.g., professional networks, cold outreach, and academic contacts) or via a sponsor-facilitated introduction used solely to reach otherwise hard-to-reach practitioner roles.

Time-to-Fill Evidence




Figure 2 Time-to-Fill Evidence by Industry. Source: Author's analysis; n = 39

Regulatory Barrier Evidence

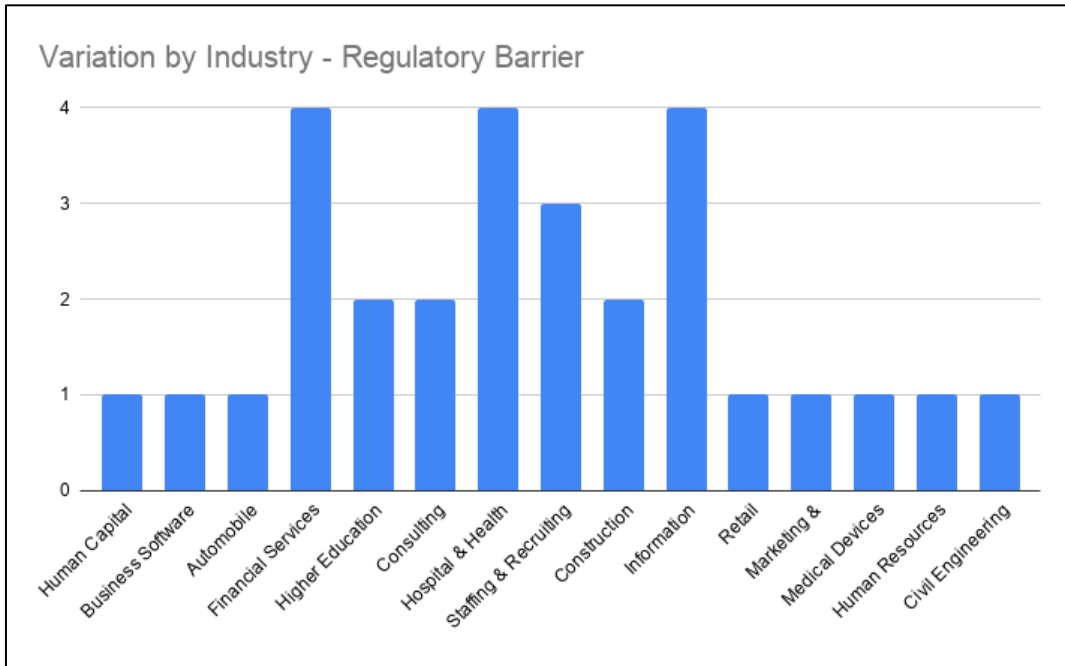


Figure 3 Regulatory Barrier Evidence by Industry. Source: Author's analysis; n = 29

Governance and Training Evidence

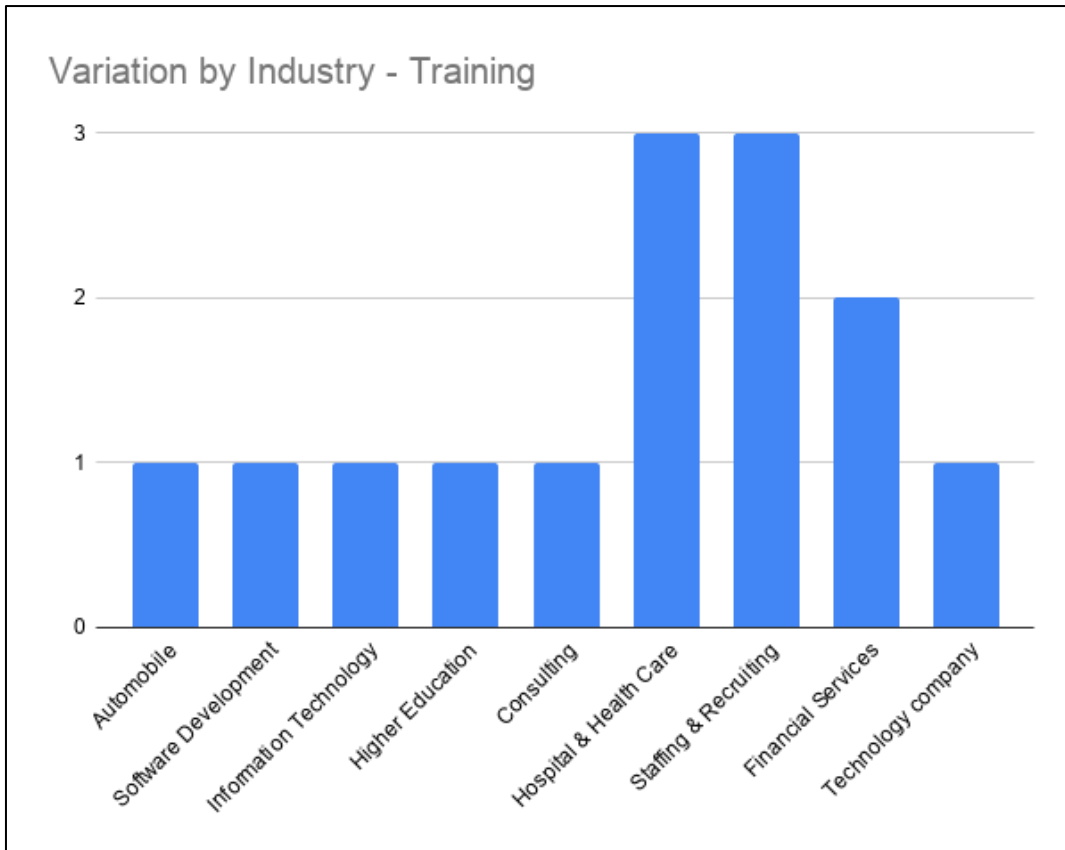


Figure 4 Governance and Training Evidence by Industry. Source: Author's analysis; n = 16

Skills-Based Hiring Evidence

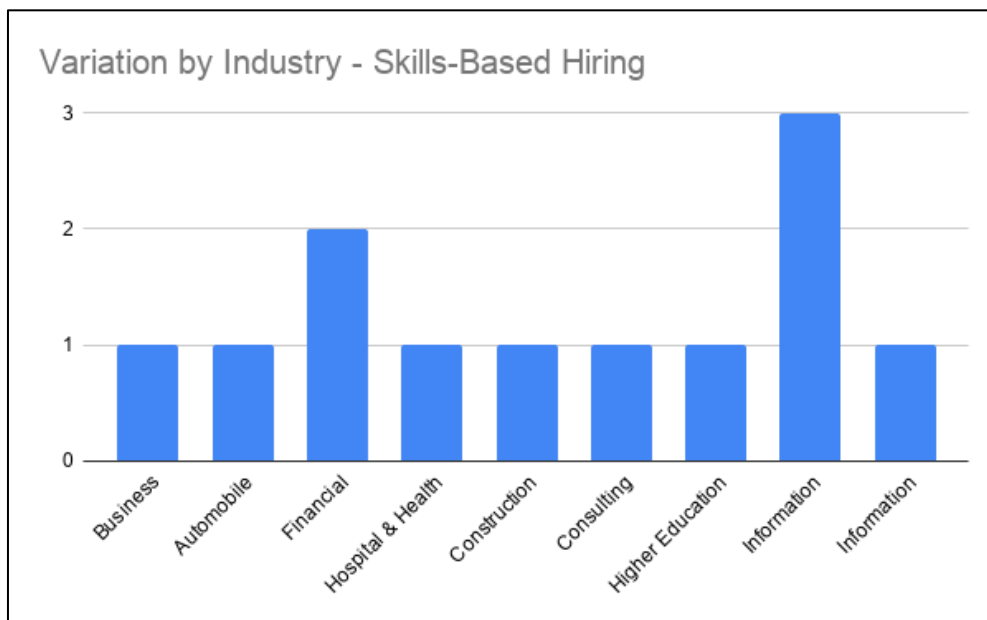


Figure 5 Skill-Based Hiring Evidence by Industry. Source: Author's analysis; n = 15

Investor Perspective Evidence

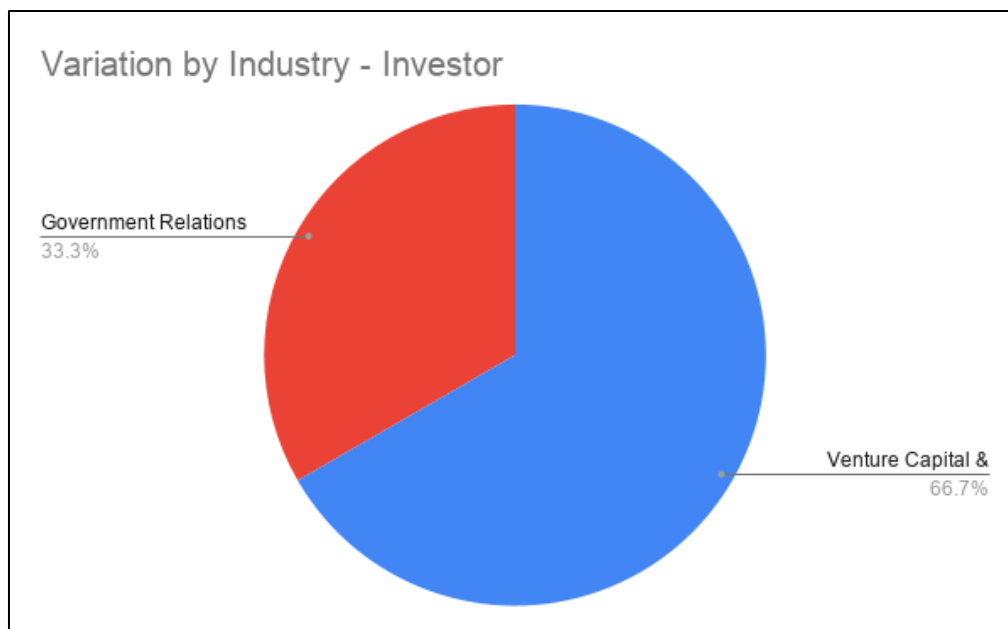


Figure 6 Investor Perspective Evidence by Industry. Source: Author's analysis; n = 3

Invalidated Hypothesis Evidence

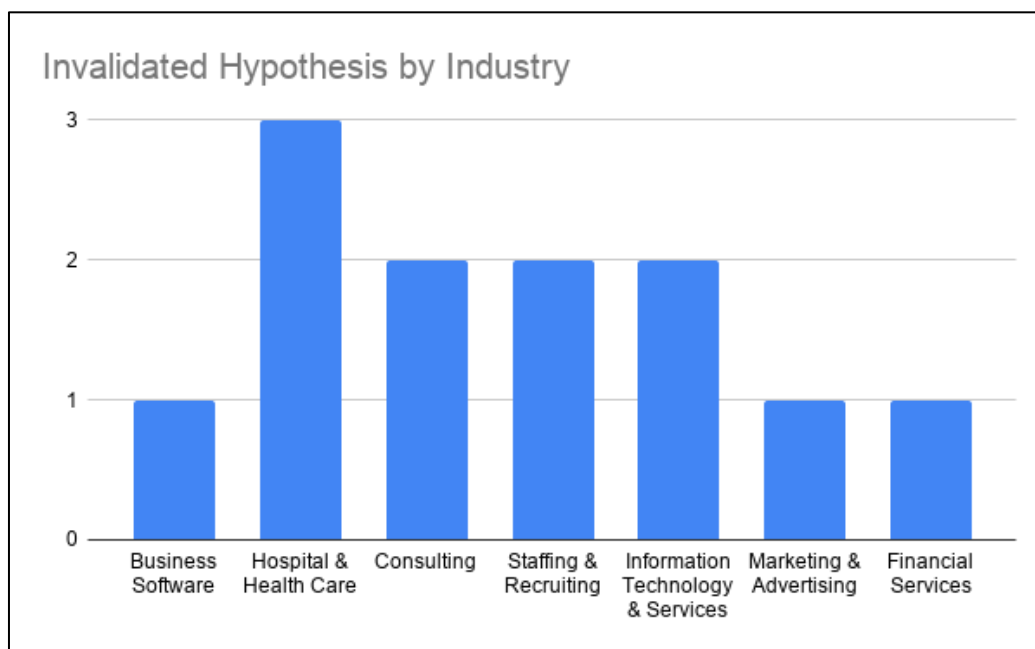


Figure 7 Invalidated Hypothesis Evidence by Industry. Source: Author's analysis; n = 12

Appendix C: AI-Assisted Transcript Analysis Documentation

C.1 Database Structure and Analytical Role of Airtable

Interview data were organized within a structured relational database in Airtable consisting of four linked tables: Contacts (participant metadata including role, industry, organization size, geography, and recruitment channel), Companies (organizational profiles linked to each contact), Hypotheses (the 15 theory-derived propositions with operational definitions and classification criteria), and Interviews (full transcripts, video recordings, and contemporaneous researcher notes linked to each participant record). This architecture allowed the researcher to query across all records systematically and to retrieve evidence organized by proposition, stakeholder type, industry, or geography.

Airtable's AI functionality (Omni) was used as a retrieval and summarization tool operating on this structured database. Its analytical role was strictly bounded: Omni surfaced candidate passages, identified recurring terminology, retrieved potential quotations, and generated descriptive summaries. All interpretation, determining whether a passage supported or contradicted a proposition, assessing evidentiary weight, and drawing conclusions, was performed by the researcher. Because the researcher personally conducted all 50 interviews, AI-retrieved outputs were always evaluated against the researcher's direct knowledge of each conversation and the contemporaneous notes logged in the same system.

Workflow Overview

The AI-assisted retrieval followed a sequential workflow:

1. Conduct and document interviews. All 50 interviews were conducted by the researcher (in-person or via video conferencing). Contemporaneous notes were recorded during and immediately after each interview and logged directly in Airtable, linked to the participant's record.
2. Transcribe and store. Video recordings were transcribed to produce full interview transcripts, which were stored in Airtable alongside researcher notes, creating a unified record for each participant.
3. Query by proposition. For each of the 15 propositions, the researcher formulated structured prompts directed at the Airtable database via Omni, retrieving hypothesis-relevant passages, recurring terminology, and candidate quotations across all 50 interview records.

4. Review and evaluate. All AI-retrieved outputs were reviewed by the researcher against the full transcript and contemporaneous notes. Each excerpt was classified as supports, contradicts, or context/unclear based on full-context reading.
5. Apply decision rules. Excerpts were retained as evidence only when they met all three retention criteria (see Section 4.5): substantive directness, corroboration across interviews, and consistency in full transcript context.
6. Summarize and count. Themes were synthesized and representative quotations selected by the researcher. Descriptive support levels (n) were computed by counting interviews with retained supporting evidence for each proposition.

Representative Prompts

Prompts were constructed around the structure of the Airtable database and directed at specific analytical needs. The following are representative examples illustrating the range of queries used. All prompts were reviewed and refined by the researcher to ensure alignment with each proposition's operational definition before outputs were evaluated.

Table 4 Representative Prompts Used to Query the Airtable Interview Database via Omni

Purpose	Representative Prompt (as directed to Airtable Omni)	Analytical Use
Geographic coverage	What are the continents of the companies of the contacts that I interviewed? I want to know the continents covered.	Verified geographic distribution across North America, Europe, and Asia for the sampling section
Keyword extraction by theme	What are the keywords for the value that the use of AI offers in talent recruitment?	Surfaced recurring language patterns across transcripts for thematic analysis; researcher reviewed all outputs
Hypothesis validation counts	Present the number of interviews completed that validate/invalidate each hypothesis. By presenting these numbers, the degree of confidence in the analysis is evident.	Generated descriptive support counts (n) per proposition; researcher verified each count against full transcript review
Interview summary and insight synthesis	Generate an executive summary for my engagements highlighting: stakeholders interviewed, learnings and findings, breakdown of the companies interviewed, and any other important insight.	Provided an initial structured overview of the dataset; researcher used as an orientation tool, not as a substitute for analytical interpretation
Interview totals and notable quotes	Present the total number of each type of interview, key interview insights, and notable quotes.	Retrieved candidate quotations for researcher review; only those consistent with full-context reading and corroborated across interviews were retained

Decision Rules for Excerpt Retention

The following decision rules were applied consistently across all propositions when evaluating AI-retrieved excerpts:

- Substantive directness: The excerpt must address the proposition's core claim, not merely reference a related theme tangentially.
- Full-context consistency: The excerpt must retain its meaning and relevance when read within the full transcript, not only in isolation.
- Corroboration across interviews: The pattern must appear independently in multiple interviews; single-instance observations were treated as insufficient for proposition validation.
- Researcher primacy: Where AI retrieval and researcher judgment diverged, including cases where the researcher's direct recall of the interview indicated a different interpretation, researcher judgment informed by first-hand interview experience and contemporaneous notes prevailed.

Note on transparency

These prompts are reproduced as actually used. They are presented here not as a formal versioned log but as a transparent record of the types of database queries employed. The analytical judgments reflected in the findings, thematic interpretation, proposition classification, and conclusion-drawing, were made entirely by the researcher and are not attributable to Airtable's Omni functionality.