



evolution_{ltd}

Research Paper

March 4, 2026

What Are Generative AI's Competitive Implications?

A Longitudental Documentation of Director Expectations

Prof. Michael G Jacobides, London Business School and Evolution Ltd, mjacobides@london.edu

M. Dalbert Ma, London Business School and Evolution Ltd, dma@london.edu

Pg	Contents
3	<u>Acknowledgments</u>
4	<u>Executive Summary</u>
5	<u>Abstract</u>
6	<u>Managerial Summary</u>
7	<u>1. Introduction</u>
9	<u>2. Theoretical Background</u>
11	<u>3. Research Design</u>
11	<u>3.1. Abductive Research</u>
11	<u>3.2. Empirical Setting</u>
11	<u>3.3. Five-Stage Design</u>
15	<u>4. Stage 1: Observations from Initial Roundtables</u>
15	<u>4.1. Methodology</u>
15	<u>4.2. Tasks and Productivity: Where Directors Agreed</u>
15	<u>4.3. Strategic Implications: Where Directors Diverged</u>
19	<u>4.4. Emergent Patterns and Transition to Stage 2</u>
20	<u>5. Stage 2: Structured Survey and Analyses</u>
20	<u>5.1. Methodology</u>
20	<u>5.2. Survey Findings</u>
24	<u>5.3. Emergent Patterns and Transition to Stage 3</u>
25	<u>6. Stage 3: Insights from Follow-Up Roundtables</u>
27	<u>6.1. Leveraging GenAI for Competitive Differentiation making</u>
27	<u>6.2. Threat of Firm Displacement</u>
29	<u>6.3. Broader Industry Disruption</u>
31	<u>7. Stage 4: Survey, Interviews, and Roundtable, 18 Months On</u>
33	<u>8. Stage 5: Interviews, Two Years On</u>
35	<u>9. Discussion</u>
35	<u>9.1. Theoretical Implications</u>
37	<u>9.2. Implications for Practice and Policy</u>
37	<u>9.3. Generalizability, Limitations, and Future Research</u>
30	<u>References</u>
42	<u>Appendix</u>
43	<u>Appendix A: Stage 1 – Open-Ended Roundtables (Feb 2024)</u>
53	<u>Appendix B: Stage 2 – Survey (Feb-Mar 2024)</u>
59	<u>Appendix C: Structured Roundtables (Apr-Jul 2024)</u>
81	<u>Appendix D: Stage 4 – Survey, One-to-One Interviews, Roundtable (Apr-Jul 2024)</u>
92	<u>Appendix E: Stage 5 – One-to-One Interviews (Jan-Mar 2026)</u>
103	<u>Appendix F: Additional Analysis – Matching Stage 2 Survey to Wayback Machines</u>
113	<u>Appendix G: Additional Analysis – Matching Stage 2 Survey to FAME Firm Data</u>

Michael G Jacobides (mjacobides@london.edu) is the Sir Donald Gordon Professor of Entrepreneurship & Innovation and Professor of Strategy at London Business School, and Evolution Ltd's Lead Advisor. He is Academic Advisor to BCG's Henderson Institute, member of the World Economic Forum's AI Governance Alliance, co-author of its White Paper on platforms and ecosystems, and ranked as one of the 50 top management thinkers. **M. Dalbert Ma** (dma@london.edu) is a PhD candidate in Strategy & Entrepreneurship at London Business School, and has worked as an Analyst in Evolution Ltd. He holds a BA in Philosophy, Politics and Economics from the University of York, an LLM in Intellectual Property and Competition Law from Erasmus University Rotterdam, and a Master in Management from IESE Business School.

The research team was supported by **Yuri Romanenkov**, Senior Advisor at Evolution Ltd and Programme Director for Next Generation Digital Strategies at London Business School, **Justinas Sukys**, Advisor at Evolution Ltd, as well as **Chinmay Bajpai**, **Netra Hirani**, **Yiru Susan Wang** (research) and **Tom Albrighton** (copy-editing).

Acknowledgments

This study was made possible by a collaboration between London Business School, the Institute of Directors in the UK (IoD), and Evolution Ltd, generously funded by the UK's Regional Innovation Fund. The original team of researchers included Yuri Romanenkov, Justinas Sukys, Chinmay Bajpai, Netra Hirani, and Yiru Susan Wang, supported by Tom Albrighton in copy-editing. Faisal Khan, Chair IoD South and Chair, IoD Science, Innovation and Technology Expert Advisory Group provided valuable guidance and encouragement, as did Roger Barker, IoD's Head of Policy. Sasha Trapani, IoD Press and Policy Officer offered stellar support and engagement, and we are indebted to the IoD members who generously offered their time, and in particular this project's Advisory Board: Pauline Norstrom, Philippe Vogeeler, Samir Chekini, Sue Milton, Yvonne Whiteley, Alistair Elder, Angus Friday, Darren Rickards, David Stringer-Lamarre, Dianne Lee, Giles Ward, Ian Clements, Perminder Ghataore, and Virginia Driver. Lord Clement-Jones, co-chair of the All-Party Political Group on AI, Lord Holmes, proposer of the Member's Bill on the UK AI Regulation Authority, and Viscount Camrose, Minister of State for AI in the UK are acknowledged for their comments on our report. General Atlantic (especially Jesse Thomas, VP leading data science) offered valuable guidance during our project on AI's impact on its portfolio that mirrored this study. We appreciate the comments and feedback from our presentation in Stanford's HAI / Digital Economy Lab, and our host Erik Brynjolfsson, as well as Kathy Eisenhardt, Angela Aristidou, and Panos Adamopoulos Moraris, and participants of two Academy of Management workshops on AI for comments and feedback. We are grateful for the detailed comments on our manuscript from our academic colleagues Ron Adner, Raj Choudhary, Donal Crilly, Kristina McElheran, Do Yoon Kim, Arianna Marchetti, Ethan Mollick, Keyvan Vakili, Gary Dushnitsky, Dan Cable, and BCG/BHI's David Zuluaga Martinez. We would also like to acknowledge Evolution Ltd and London Business School's Knowledge Exchange Fund for additional support.

Executive Summary

- Generative AI (GenAI) has provoked both hype and intense debate in the business world. Many leading consultants, academics, and industry experts are heralding it as a transformative force for both organizations and entire sectors, with some even projecting trillions in potential economic impact. A minority of academics, however, argue that the impact will be more modest in economic terms and will depend on complementary investments.
- Whatever the size or speed of GenAI's impact, the question remains: how disruptive will it truly be? Most research predicting major change has focused on simple, modular, individual tasks—yet such capabilities may not necessarily aggregate to organizational or sectoral levels, especially when we consider the historical precedent of earlier technological shifts.
- Our comprehensive study gathered valuable insights into strategic decision-making and resource allocation in the context of GenAI adoption. We employed a three-stage approach: initial qualitative roundtables with 96 executives, a quantitative survey of 217 executives, and a final qualitative stage involving 163 executives.
- Our findings reveal significant disparities in expectations of disruption, both between and within sectors, suggesting the potential for heterogeneous effects. We also distinguish between GenAI-driven displacement, i.e. the wholesale challenge to value-add activities of firms and entire sectors, and its impact on supporting differentiation for firms leveraging GenAI. Some settings more exposed to GenAI will see both occur simultaneously.
- We found that industry-specific features play a key role in how executives view the threat of displacement from GenAI. For example, executives in heavily regulated industries or those requiring high output accuracy feel more protected. The forces driving displacement fears and differentiation opportunities sometimes overlap, suggesting that in certain sectors, GenAI might benefit some firms while challenging others.
- At the organizational level, executives see assets such as proprietary data and tacit knowledge as crucial for leveraging GenAI for competitive advantage, while organizational structure, particularly modularity, plays a key enabling role in deployment.
- Our findings suggest that organizations should develop capabilities in proprietary data, data cleanliness, organizational flexibility, and tacit knowledge to leverage GenAI effectively. Strategies should balance sector-specific dynamics with unique organizational capabilities.
- Policymakers need to consider frameworks for data democratization, sector-specific interventions, and broader analytical approaches to competition. These efforts should aim to promote equitable GenAI integration while addressing its wide-ranging implications on market dynamics and cross-sector disruption.
- Our research underscores the need to look beyond the hype surrounding GenAI and consider broader organizational and sectoral dynamics. This more nuanced approach reveals that GenAI's impact is likely to be uneven and context-dependent, challenging simplistic narratives of universal disruption or transformation.
- Our paper offers a rigorous basis for helping companies rethink their approach to GenAI, with more focused frameworks / tools for guiding strategy being developed by Evolution Ltd.

Abstract

How do senior decision-makers assess the competitive implications of generative AI (GenAI)? Using an abductive, multi-stage study of 531 UK corporate directors (exploratory roundtables, survey, matched-stimulus roundtables, and follow-up interviews, 2024–2026), we document that directors distinguish three dimensions of GenAI's competitive implications: competitive differentiation, firm displacement, and industry disruption. Expectations vary markedly even within the same sub-sector. This heterogeneity is patterned by how directors frame their firms' value propositions—especially perceived modularity and reliance on pattern recognition—and by complements and buffers such as proprietary data, relationship-intensive delivery, and regulatory/liability constraints. Across waves, directors' competitive expectations remain stable, consistent with durable interpretive frames during early diffusion.

Managerial Summary

Directors do not form one “overall” view of GenAI. They separate three questions: will it help us differentiate, make us vulnerable to displacement, or reshape the sector through industry disruption? Even within the same sub-sector, leaders often disagree—because they frame the value proposition differently. Two lenses matter: whether the offering is modular and whether it is largely pattern-recognition-based. Perceived complements and buffers—proprietary data, relationship-intensive delivery, and regulatory/liability requirements—shape where GenAI is seen as amplifying advantage versus commoditizing work. While views of capabilities evolve, competitive expectations are comparatively stable, so making assumptions explicit and stress-testing them may be as important as tracking the technology

Keywords: artificial intelligence, generative ai, strategic decision-making, competitive advantage, technological disruption, pattern recognition, modularity

1. Introduction

A growing body of research documents that generative AI (GenAI) can perform a wide range of knowledge-work tasks—writing, ideation, analytical synthesis, coding—and that the resulting productivity gains can be significant (Appel *et al.*, 2025; Daniotti *et al.*, 2026; Dell’Acqua *et al.*, 2023; Noy and Zhang, 2023). Yet knowing what a technology can do does not directly tell us what it will mean for firms and sectors. Strategy research has long recognized that firms’ existing assets and competitive positions shape how decision-makers interpret a new technology’s implications, and that these interpretations shape how a technology’s capabilities influence resource-allocation decisions and, ultimately, competitive trajectories (Eggers and Kaplan, 2009; Eggers and Park, 2018; Wu, Wan, and Levinthal, 2014). These theoretical insights suggest that similar dynamics should characterize firms’ responses to GenAI—yet systematic evidence remains scarce. We lack empirical understanding of how senior decision-makers assess GenAI’s competitive implications, and what drives heterogeneity in those assessments.

In this paper, we take a first systematic step toward understanding how senior decision-makers assess GenAI’s competitive implications, by examining how corporate directors interpret what the technology means for their firms and sectors. GenAI presents a distinctive challenge for competitive assessment: because the technology applies broadly across task types, functions, and industries (Bick, Blandin, and Deming, 2026; Eloundou *et al.*, 2024; Felten, Raj, and Seamans, 2023; Hartley *et al.*, 2024), the connection between its capabilities and its competitive implications is far from self-evident (Csaszar and Jia, 2026; Gans, 2025). Given this uncertainty, directors’ interpretations of *what* matters and *why* serve as precursors to firm decisions on resource allocation and competitive positioning (Eggers and Kaplan, 2009; Tripsas and Gavetti, 2000), offering an possible window into how firms and industries are likely to respond to GenAI.

Empirically, we draw on an abductive, mixed-methods study (Mitchell *et al.*, 2022; Timmermans and Tavory, 2012) of 531 corporate directors, primarily from the UK across five stages between February 2024 and March 2026. By returning to the same population of directors over two years—through roundtables, structured surveys, matched-peer sessions, and follow-up interviews—we were able to surface, quantify, and probe the mechanisms underlying directors’ competitive assessments, and to examine whether those assessments proved stable as the technology evolved. To triangulate beyond the IoD population, we also conducted roundtables with external, predominantly U.S.-based directors throughout the stages of our study and supplemented our core findings with archival textual analysis and matched firm-level data.

We document four sets of findings. First, directors conceptualize GenAI’s competitive implications along three dimensions. (1) *Competitive differentiation* concerns whether GenAI enables an incumbent to improve or reconfigure its value proposition relative to rivals. (2) *Firm displacement* concerns whether GenAI enables imitation, disintermediation, or capability substitution that threatens incumbents’ viability. (3) *Industry disruption* concerns whether GenAI reorganizes value creation and capture at the sector level, altering roles, boundaries, and the architecture of competition. Across these dimensions, directors exhibit substantial

heterogeneity in expectations—even within sub-sectors where they agree on GenAI’s task-level capabilities.

Second, this heterogeneity does not correspond to task- or occupation-level exposure, but to how directors construe their firms’ value propositions and complementary assets in areas such as proprietary data, risk appetite, relationship-intensive delivery, and regulatory intermediation. These complements appear salient not only because they shape what firms can do with GenAI, but because they structure what directors anticipate that competitors and entrants can credibly replicate. Third, two organizational perceptions consistently structure expectations: the extent to which a firm’s value proposition is perceived as *modular* and the extent to which it is perceived as reducible to *pattern recognition*. Directors within the same sub-sector diverge sharply on these perceptions, describing ostensibly similar work in markedly different terms. Textual analysis matching director responses to archived website data suggests that objective positional differences account for only part of this variation. Fourth, directors’ competitive expectations remain remarkably stable across the two-year period of our documentation, even as directors substantially revise their assessments of GenAI’s task-level capabilities upward—consistent with durable interpretive frames.

Our contribution is to provide empirical texture to a relationship that, to our knowledge, has yet to be systematically explored: that between GenAI’s widely discussed capabilities and firms’ competitive expectations. Research has moved quickly to document what GenAI can do; what is largely missing is a structured account of how senior decision-makers connect those capabilities to expectations about competition, and why those connections vary so markedly even among directors who agree on what the technology can do. Our findings reveal rhyme and reason to this heterogeneity and suggest a way to begin unpacking it, providing an empirically grounded map of how decision-makers translate GenAI’s capabilities into the competitive expectations that guide resource allocation and strategic positioning, and generating propositions to guide future research connecting task-level technological change to firm-level competitive response and outcomes.

2. Theoretical Background

Research on technological change has long established that the competitive implications of a new technology cannot be read directly from its capabilities (see Eggers and Park (2018) for a review). Between what a technology can do and what it means for firms and industries lies a critical interpretive step: firm decision-makers must assess which of the technology's capabilities are operationally and competitively relevant; how they interact with existing assets and market positions; and whether they enable, threaten, or restructure the basis of competition (Dutton and Jackson, 1987; Henderson and Clark, 1990; Jackson and Dutton, 1988; Kaplan and Tripsas, 2008). Because this assessment is shaped by firms' existing positions and heterogeneous prior experiences, the same technological possibilities give rise to markedly different expectations—and different strategic responses—across firms (Shane, 2000; Tripsas and Gavetti, 2000; Wu *et al.*, 2014). Understanding the competitive implications of any new technology thus requires bridging two bodies of evidence: what the technology can do and how decision-makers connect those capabilities to form expectations about competition for their firms and sectors.

For GenAI, the first of these is increasingly well served: the speed and breadth of GenAI's diffusion has generated a substantial body of research documenting what the technology can do. One stream examines the locus and magnitude of GenAI's productivity effects. Beyond macro-economic models that project economy-wide implications (e.g., Acemoglu, 2024), a number of carefully designed studies measure productivity gains within particular domains, generally at the level of well-defined tasks. These studies document significant benefits in areas such as writing (Noy and Zhang, 2023), strategic decision-making (Doshi *et al.*, 2024), consulting problem-solving (Dell'Acqua *et al.*, 2023), customer service (Brynjolfsson, Li, and Raymond, 2025b), and coding (Hoffmann *et al.*, 2025; Peng *et al.*, 2023). Complementary work surveys the increasing range and complexity of tasks GenAI is being used for (Chatterji *et al.*, 2025; Handa *et al.*, 2025; Makridis, 2025). A second stream assesses the breadth of GenAI's potential reach across the occupational landscape. Drawing on Brynjolfsson, Mitchell, and Rock (2018), this work decomposes occupations into constituent tasks and evaluates each task's susceptibility to AI automation (see Eloundou *et al.*, 2024; Svanberg *et al.*, 2024; Tomlinson *et al.*, 2025). Related research examines GenAI's implications for labor markets, integrating the displacement effects of automation with the augmentation effects of human–AI collaboration¹ (Brynjolfsson, Chandar, and Chen, 2025a; Felten *et al.*, 2023; Johnston and Makridis, 2025; Manning and Aguirre, 2026).

The second body of evidence—concerning how decision-makers connect these capabilities to competitive expectations—is far thinner. Practitioner accounts have asserted that GenAI will reshape competition (Choudary, 2025; Suleyman and Bhaskar, 2023), and consulting reports have intimated expectations of change among business leaders (BCG *et al.*, 2024), but

¹ Some work points out that such improvements may not translate into aggregate productivity benefits at the economy-wide level, since this would require entire processes to change and organizations to alter their means of producing (Martinez, 2025). However, the consensus is that productivity benefits will be substantial, if unevenly distributed, with a 'jagged frontier' whereby some activities will benefit from GenAI and others not (Dell'Acqua *et al.*, 2023; Iansiti and Lakhani, 2020; Jacobides, Brusoni, and Candelon, 2021).

academic work addressing the question directly remains limited. Jacobides et al. (2021), writing pre-GenAI, consider how AI more broadly reconfigures industry architecture: who controls key bottlenecks, how complementarities shift, and how firms can shape their sectors and occupy advantaged positions. Iansiti and Lakhani (2020) examine how digital infrastructure can create competitive advantage within the firm, and suggest these findings extend beyond their sample of technology pioneers. Recent work has begun to examine strategic repositioning in response to GenAI, though at the individual rather than firm level: Yiu et al. (2025) document how freelancers on the freelancing platform Upwork adjusted their specializations and repositioned following ChatGPT's launch. These contributions provide relevant foundations, but systematic evidence on how firm-level decision-makers connect GenAI's capabilities to expectations about competition remains scarce—leaving open the question of whether and why those expectations vary.

The broader literature on technological change, however, provides substantial foundations for how we may unpack this question. Central to this literature is the insight that firms' existing assets and market positions play a dual role in technological transitions, functioning as both 'pipes' and 'prisms'—that is, as resources that shape what firms can credibly do with a new technology and as interpretive lenses that direct attention toward certain competitive trajectories over others (Kaplan and Tripsas, 2008; Teece, 1986; Wu *et al.*, 2014). Broadly applicable technologies are also known to require complementary investments—in data, human capital, and organizational redesign—to realize their competitive potential (Agrawal, Gans, and Goldfarb, 2021; Baldwin, 2024; Bresnahan, 2024; Bresnahan and Trajtenberg, 1995; McAfee, 2024; Puranam, 2025), though which complements prove decisive depends in part on how decision-makers interpret what the technology demands. Understanding which organizational attributes decision-makers perceive as relevant to a new technology—as both resources and interpretive frames—thus provides insight into the competitive trajectories that firms are likely to pursue (Benner and Waldfoegel, 2016).

For a technology whose competitive consequences have not yet fully materialized, examining how decision-makers assess those consequences offers the most direct route into the question. Prior research has established that executive perceptions of new technologies are not merely reflections of an underlying reality awaiting study, but active inputs that shape firm responses and, collectively, the technological trajectories that unfold (Benner and Tushman, 2002; Eggers and Kaplan, 2009; Tripsas and Gavetti, 2000). How executives frame a technology—as reinforcing or disrupting existing advantages—guides resource allocation, strategic positioning, and the pace and direction of technology engagement (Dutton and Jackson, 1987; Gilbert, 2005, 2006; Jackson and Dutton, 1988). Examining what corporate directors attend to when assessing GenAI's competitive implications thus serves as a precursor to studying actual competitive outcomes once they unfold, offering a first empirical look at a question that has so far been engaged only speculatively.

3. Research Design

3.1 Abductive Research

Our research employs an abductive approach—one that begins not with hypothesis testing but with sustained engagement with a puzzling phenomenon, allowing theoretical insights to emerge through iterative dialogue between data and existing frameworks (Mitchell *et al.*, 2022; Timmermans and Tavory, 2012). GenAI presents precisely the conditions that favor abduction: a phenomenon generating considerable speculation but little systematic evidence (Choudhury, Moeen, and Wormald, 2025). Our design reflects this logic: observations from each stage informed the design of the next, with each successive stage addressing questions that emerged from prior findings until theoretical saturation was reached (Timmermans and Tavory, 2012).

3.2 Empirical Setting

We partnered with the Institute of Directors (IoD) through a collaboration funded by the UK's Regional Innovation Fund. Established in 1903 and incorporated by Royal Charter, the IoD is a paid professional membership organization for company directors and senior business leaders across the private, public, and third sectors that serves as a platform for members to engage with public policy, technological developments, and governance issues through events, networking, and advocacy to government. This setting controls for executive seniority and geographic context while providing broad cross-sectoral variation. The IoD's substantial representation from professional services—sectors ranking high on task and occupational exposure measures (Eloundou *et al.*, 2024; Felten *et al.*, 2023; Svanberg *et al.*, 2024)—enabled early examination of granular within-sector differences rather than only cross-industry comparisons.²

3.3 Five-Stage Design

Our investigation unfolded across five stages between February 2024 and March 2026. **Figure 1** summarizes the design; detailed procedural specifications appear in the Appendices A-E, with stage-specific methodological choices described in each findings section. **Table 1** provides the corresponding participant composition across stages (session types, counts, and sector coverage). To triangulate beyond the IoD population, we also conducted external roundtables with non-IoD, predominantly US-based directors at Stages 3, 4, and 5.³ Additional analyses matching survey responses to website data and to firm-level financial data are detailed in Appendices F and G.

² Two sources of self-selection are inherent in this setting: the IoD's membership base is itself self-selected, as directors who join a professional organization may differ from the broader executive population; and participation in our research was voluntary, with the multi-stage design likely compounding selection toward directors with sustained interest in GenAI. These are inherent limitations of perceptual research involving senior decision-makers, who cannot be randomly sampled or compelled to participate. However, because our inquiry centers on heterogeneity in expectations and the organizational attributes that drive it—rather than on estimating population-level magnitudes—self-selection into engagement is less threatening to our inferences. We bound our findings accordingly: they speak most directly to executives actively engaging with GenAI rather than to the broader population of corporate directors.

³ Findings from external roundtables corroborated the core patterns documented here. We report these, along with additional observations, in the Appendix to maintain focus on the IoD population that constitutes our primary sample.

Figure 1: Progression of the abductive analysis




Table 1: Participant Composition Across Stages

Stage	Period	Session / Method	N	Sector Coverage
1	Feb 2024	Non-regulated professional services	19	Management consulting, technology consulting, engineering and construction consulting, healthcare and life sciences consulting, education consulting
		Omnibus (1)	20	Education, management consulting, finance and risk management, technology and energy, data analytics, professional training, software development
		Media & communications	19	Communications, PR and marketing, media production, advertising, market research, business services
		Regulated professional services	16	Legal services (corporate law, litigation, patent and trademark law, immigration), insurance, banking, management consulting
		Omnibus (2)	22	Management consulting, technology consulting, financial services, banking, engineering and construction consulting, pharma, life sciences
2	Feb–Mar 2024	Structured survey	221	Professional services, information & communication, financial services, education, manufacturing, construction, and 9 additional sectors†
3	Apr–Jul 2024	High pattern-recognition importance	23	Pharmaceuticals, biotechnology, cybersecurity, insurance, management consulting, ESG assessments, environmental planning and landscape design
		Low pattern-recognition importance	13	Management consulting, venture capital, software development, environmental services, alternative dispute resolution, AI ethics certification, digital agency
		High proprietary data importance	17	Pharmaceuticals, management consulting, legal services (corporate, litigation), insurance, banking, technology advisory, corporate restructuring
		Low proprietary data importance	15	Telecommunications, management consulting, venture capital, chemical manufacturing, marketing, medical services, construction, education
		High competitive uncertainty	11	Management consulting, digital banking, technology solutions, risk and compliance advisory, real estate advisory, manufacturing
		Low competitive uncertainty	18	Management consulting, technology, venture capital, legal services, chemical manufacturing, manufacturing, meteorology, food and beverage export
		Modularity (mixed)	15	Healthcare, pharmaceuticals, financial services, technology, management consulting, legal services, construction equipment supply, language translation
		External roundtable (1)	21	Management consulting, banking, insurance, legal services, pharmaceuticals, telecommunications, ESG assessment, education
		External roundtable (2)	30	Private equity, venture capital, consulting, legal services, tech, banking, retail
4	Apr–Jul 2025	Follow-up survey	41	Consulting, financial services, technology, policy, construction, pharmaceutical advisory, and 7 additional sectors†
		One-to-one interviews	22	Software, management consulting, hardware, technology consulting, construction & architecture, education, edtech, venture capital, copywriting, law, financial services, accounting, insurance
		External roundtable	27	Management consulting, financial services (banking, insurance), venture capital, healthcare, media, retail & consumer goods, logistics & postal services, IT services, AI research, executive search, policy, academia
5	Jan–Mar 2026	One-to-one interviews	41	Consulting, financial services, technology, cybersecurity, construction & architecture, non-profit, policy, insurance, pharmaceutical advisory, edtech
		External roundtable	23	Private equity, management consulting, financial services, technology, policy
		External roundtable	24	Private equity, venture capital, management consulting, financial services, technology, policy, software, legal services, pharmaceutical advisory

Notes: Total unique participants across all stages: 531. IoD = Institute of Directors. Stage 1: 102 registered, 96 attended. Stage 2: 486 total responses, 221 fully completed; ~11% response rate from IoD UK membership. Stage 3 paired roundtables stratified participants by responses from Stage 2; external roundtables comprised participants outside the IoD sample. Stage 4 survey included 27 matched Stage 2 respondents (within-person panel). Stage 4 external roundtable: 30 invited, 27 attended. † Full sector distribution in corresponding Appendix.

Stage 1 (February 2024): Open-ended roundtables. We conducted five roundtables with 96 directors: separate sessions for regulated professional services, non-regulated professional services, and media & communications, alongside two omnibus sessions from a wider array of sectors. The timing—15 months after the launch of ChatGPT—meant uncertainty remained high and directors had made few public commitments that might constrain candid assessment. We chose roundtables deliberately: we sought to observe where directors naturally agreed and disagreed in real time, and sector-based groupings enabled us to examine both within-sector and cross-sector patterns.

Stage 2 (February–March 2024): Structured survey. Patterns from Stage 1 required systematic quantification. A structured survey also offered a methodological complement: where roundtables captured how directors articulated and refined positions in real time, individual responses captured private assessments unconstrained by group dynamics. We administered a pre-tested instrument to 221 directors, enabling triangulation between roundtable contributions and private responses, and partitioning directors into theoretically meaningful groups for Stage 3.

Stage 3 (March–May 2024): Structured roundtables with matched peers. To probe mechanisms underlying earlier patterns, we designed roundtables inverting the logic of Stage 1. Rather than grouping by sector, we grouped by response pattern—matching directors who shared similar assessments on some dimensions but diverged on others. Each roundtable exposed matched peers to identical GenAI scenarios, enabling systematic comparison of interpretive frameworks.

Stage 4 (April–July 2025): Follow-up survey, interviews, and external roundtable. Over a year separated our initial data collection from Stage 4—a period during which GenAI capabilities expanded and directors had opportunities to form, and publicly commit to, competitive positions. Directors who had made such commitments might present consistent narratives in group settings regardless of private reassessment. We shifted to individual interviews and a re-administered survey, examining whether earlier patterns reflected durable frames or transient reactions.

Stage 5 (January–March 2026): Follow-up interviews. These follow-up interviews were conducted to contextualize and deepen insights from Stage 4, focusing on how directors reflected on the persistence or evolution of their competitive assessments over time.

The progression across stages reflects deliberate methodological choices to ensure qualitative rigor. Early-stage roundtables leveraged group dynamics to surface disagreement; later-stage individual methods mitigated those same dynamics when posturing might otherwise dominate. Sector-based groupings in Stage 1 revealed initial patterns; response-based groupings in Stage 3 isolated mechanisms. The extended temporal window between Stages 3, 4, and 5 enabled us to observe whether conceptions remained stable as the technology evolved.

4. Stage 1: Observations from Initial Roundtables

4.1 Methodology

We began our abductive investigation with a focal question: *How do directors perceive GenAI's implications for their firms and sectors?* We framed this broadly, allowing directors to surface the dimensions they considered most salient rather than imposing categories *ex ante*.

Interviewers asked roundtable participants how they saw GenAI within their organizational context, probing experiences, expectations, and the potential for transformation. Participants were encouraged to share perspectives, respond to and challenge others, and keep their answers concrete and context-specific. Detailed procedures and supplementary observations appear in Appendix A.

4.2 Tasks and Productivity: Where Directors Agreed

Across all five roundtables, directors demonstrated strong consensus on GenAI's task-level capabilities. In fact, GenAI's task capabilities and productivity implications consistently emerged as the first topic. Partners at corporate law firms acknowledged its effectiveness in contract review, due diligence, and legal research. Directors at M&A advisory firms and investment banks agreed on its utility for financial modeling and pitch materials. Strategy consultants cited its capacity for desk research, framework generation, and first-draft deliverables. PR and communications directors concurred on its capabilities in generating copy variations and emulating brand voice.

Directors discussed these capabilities in productivity terms—efficiency gains, time savings, cost reductions, reduction in workforce. A director at a software firm described how an AI coding assistant “saved [a senior developer] two weeks of work.” A partner at a regional law firm observed that contract review tasks previously requiring junior associates could now be completed “in a fraction of the time.” When facilitators asked whether directors disagreed with colleagues' task-level assessments, substantive disagreement was rare. To the extent debate arose, it centered on the degree to which GenAI could perform or assist with a given task—not on task breadth itself.

4.3 Strategic Implications: Where Directors Diverged

Where productivity discussions yielded consensus, discussions of strategic implications surfaced heterogeneity not only across sectors but within sub-sectors—even among directors whose task-level assessments were closely aligned. **Table 2** presents representative perspectives. We focus here on four sub-sectors where matched comparisons proved particularly revealing.

Table 2: Enablers and barriers linked to GenAI experience (Stage 1)

Enablers	Barriers
<i>Proprietary data being supercharged through GenAI</i> “We’ve found that we’re essentially sitting on a goldmine. We have this library of press releases that we’ve not been able to properly use until now. Now we can replicate at scale our writing style.” (Director in PR firm, Roundtable 3) “What we are doing currently with GenAI is using our library of prior data and super-scaling it with the tech, [...] seeing if that can even create new business opportunities for the firm.” (Director in law firm, Roundtable 4) <i>Tacit knowledge in human labor complementing GenAI’s lack of judgement</i> “The technology is impressive, but it’s not going to take your job, because it can’t. It doesn’t have human intelligence. But in the right human hands, it’s very powerful.” (Managing Director in construction consultancy, Roundtable 1) “ChatGPT can give me a bunch of facts, but [...] it’s still the job of the consultant, and his experienced judgement, that is needed to determine which facts matter and which don’t.” (Director in strategy consultancy, Roundtable 1) <i>Technical expertise in humans correcting and verifying GenAI’s output inaccuracy</i> “Normally, we end up with a team of 20 or 30 developers writing code [...] There will almost always be a bug [...]. Now, our company got an AI tool [...] It still required [...] a high-level developer to check that code [...] But he said it saved him two weeks of work.” (Director in software development firm, Roundtable 2) <i>Relational capital and client trust buffering the firm against selection issues</i> “It’s a matter of trust about whether the quality of your advice is good enough. If you’re using these AI tools as co-thinkers, [...] can you prove to me that what you’ve got through [...] AI is trustworthy?” (Logistics director, Roundtable 1) <i>Organizational risk appetite enabling adoption and first-mover advantage</i> “If you’re an organization that sits in the pie of a market, you just have to eat your little slice, you’re finished. Because someone else who isn’t will kill you at a quicker pace than anything normally happens in the commercial market. If you’re someone who’s willing [...], you don’t want that pie, you want to create a pie.” (Director in educational consultancy, Roundtable 2)	<i>Lower modularity increasing the complexity of deploying the technology</i> “There is just a lot of layers to get through before we can use the technology. GenAI is good at replacing [individual] tasks, but an organization is a system of tasks.” (Director of management consulting firm, Roundtable 1) “But the problem is big corporations are tied to standards [for] standardization. The deployment of new technologies, things like that, are not very favorable [...] In a small organization, it’s very welcome. But a bigger corporation would see it as risk.” (Director of education technology firm, Roundtable 2) “A few years ago, we acquired this small tech firm in order to digitalize our backend. It was a complete failure; you realize that the technology stacks just don’t fit together. The engineers [...] spoke different languages. It’s a similar experience for us with GenAI.” (Director of investment bank, Roundtable 5) <i>Organizational leadership being uninterested or not understanding the technology</i> “CEOs are at the position they are at because what they have done has worked for the past decade. It’s not in their interest to invest in this transformative and uncertain technology.” (Director of consulting firm, Roundtable 4) “Executives are, at the end of the day, risk-averse people. It does not help that the ROI of this technology is difficult to measure or understand.” (Director of financial services firm, Roundtable 2) <i>Inability to identify and/or attract the right talent to deploy the technology</i> “It’s not that we don’t want to use the technology, and we can clearly see some of our competitors already using it. We just don’t really have the right people [...] to lead its implementation.” (Director at advertising firm, Roundtable 3) “Our organization is small and not as well-known as some of these bigger players in our sector. We struggle with the pull to bring [in] the right engineering talent, or even to know who they are.” (Director at EdTech firm, Roundtable 2) “We’re already experiencing what was a talent drain in a few years before GenAI and over COVID. Lots of people leaving the local sector to go to other sectors, and a lot of technical-minded people [...] This means that only the top companies with the big budget and big resources are now able to attract people back.” (Director at technology firm, Roundtable 2)

Roundtable 1: Strategy Consulting (Non-Regulated Professional Services)

Directors at strategy consulting firms agreed that GenAI could perform many analytical tasks: synthesizing research, generating frameworks, structuring problems, and proof-checking. Yet when discussion turned to strategic firm and sector-level implications, disagreement surfaced.

A managing director at a mid-sized strategy boutique (approximately 200 consultants, FTSE 250 clients) described GenAI as prompting fundamental reexamination: “We are actively reevaluating our corporate purpose [...] searching for new sources of value creation.” He described internal conversations about whether traditional deliverables—strategy decks, market analyses, competitive assessments—would retain value when clients could generate such outputs themselves.

A partner at a comparably sized consultancy pushed back, questioning whether analytical outputs had ever truly embodied value: “A consulting project is not just advisory [...] it’s mainly implementing change in your company. And change management is done by people.” The firm’s value, he argued, lay in relational and implementation work—navigating organizational politics, building client commitment, managing transformation dynamics. When the first director pressed on whether implementation was truly defensible, the second countered that value lay in bundling: “They come to us because we do both. Unbundle it and you’re competing with [technology consultancy] on one side and boutique facilitators on the other.”

A senior partner at a global firm’s London office offered yet another framing. He saw opportunity in combining GenAI with “decades of proprietary frameworks and client data” to create new forms of insight delivery, suggesting this would shift value capture within the sub-sector. This exchange illustrated a recurring pattern: directors in the same sub-sector, with aligned task-level assessments, diverging on which features of their value proposition mattered most.

Roundtable 2: Legal Services (Regulated Professional Services)

Among directors at corporate law firms, task-level consensus was equally strong—GenAI performed impressively on document review, contract analysis, legal research, and first-draft preparation. Competitive implications prompted divergence on firm-level positioning, but notable consensus on sector-level stability.

A partner at a mid-market firm characterized GenAI as an efficiency tool that would not fundamentally alter the profession’s structure, pointing to professional liability, regulatory oversight, and certification: “You can’t sue the AI if it gives you wrong advice. Clients need someone who carries professional indemnity insurance, someone the SRA regulates, someone they can hold accountable.” This sentiment was met with notable agreement across the roundtable.

Another partner at a mid-market firm was less sanguine about his firm’s positioning. He acknowledged regulatory protections but worried that GenAI might enable clients with sophisticated in-house teams to handle more matters internally: “The magic circle will be fine—their clients need the brand, the cross-border capability, the relationship partners. The question is whether there’s still a middle market in five years, or whether it hollows out.”

A senior director at a legal technology company suggested that both partners were underestimating the potential for structural change in service delivery, pointing to alternative legal service providers as evidence that the “traditional partnership model” was already under pressure. Yet, when she questioned whether regulatory barriers would hold, another partner pushed back firmly: “Parliament isn’t going to change the Legal Services Act because of ChatGPT.” The mid-market partner was more circumspect but agreed wholesale disruption was unlikely. Directors expected sector boundaries to remain intact, though differed over how advantage might be redistributed.

Roundtable 3: Financial Services (Omnibus)

A similar pattern emerged among directors in M&A advisory and corporate finance. Directors agreed GenAI excelled at analytical components—comparable analyses, due diligence synthesis, transaction modeling—while diverging on firm and sector-level implications.

A managing director at a specialist M&A boutique worried about “disintermediation at the margins”—not dramatic displacement, but gradual fee compression as analytical work became more transparent to clients. A director at a bulge-bracket bank countered that “the analysis was never really the point”—her firm’s value resided in relationships, market access, and coordinating negotiations under time pressure. “No one hires us for our slides.”

Yet, directors all agreed that the advisory sector itself would not be fundamentally upended. A director at a wealth advisory firm captured the prevailing view: “At the end of the day, someone needs to sign off. Someone needs to be accountable. Clients aren’t going to rely on an algorithm for a decision this consequential.” There was strong agreement here. Requirements for human accountability and professional liability appeared to anchor expectations about sectoral stability even as directors diverged on competitive dynamics within that structure.

Roundtable 4: Public Relations (Media and Communications)

The PR and communications roundtable surfaced a different pattern—more varied responses and, for some directors, more acute concern than in regulated professional services.

A director at a mid-sized integrated communications agency expressed profound concern: “If a client decides their marketing team can do this with ChatGPT, there’s nothing structurally stopping them.” A director at a specialist brand consultancy pushed back by emphasizing accumulated understanding of specific clients’ brand identities: “We’re essentially sitting on a goldmine. We have this library of materials [...] that we’ve never been able to properly use until now.”

A third director—at a crisis communications firm—argued his value proposition was largely insulated: “In our slice of the market, clients and readers can tell AI-produced content. We have a specific style of writing that takes years for us to teach our juniors.” The deliberation that followed was more anxious than in regulated sectors. When the crisis director suggested “human-produced” work was defensible, the integrated agency director countered: “That’s fine for crisis work, but most of our revenue is content production. Most clients most of the time are buying volume, not judgment.” Directors in less regulated sub-sectors spoke about more

actively searching for firm-specific sources of defensibility and exhibited greater heterogeneity in what they identified.

4.4 Emergent Patterns and Transition to Stage 2

Our Stage 1 roundtables surfaced several patterns that shaped subsequent inquiry.

Three dimensions of competitive implications. A consistent structure emerged in how discussions unfolded. Directors engaged intensively on particular topics—debating until apparent saturation—before distinct topics surfaced renewed deliberation. Across roundtables, these discussions clustered into three dimensions that directors discussed as broadly separate phenomena: the ability to leverage GenAI for competitive differentiation, the threat of firm displacement, and broader sectoral disruption. Notably, directors invoked the term “disruption” most frequently for sector-level transformation, reserving the term for industry restructuring rather than firm-level dynamics.

Decoupling of productivity and competitive discussions. Directors naturally gravitated toward productivity discussions first, describing use cases and efficiency gains with considerable consensus. Where disagreement arose, it typically concerned trajectory—how quickly capabilities would improve or whether current limitations would persist. Yet, these exchanges led primarily to operational concerns: hiring when entry-level tasks are automated, developing tacit knowledge when juniors rely on AI, structuring training pipelines. When facilitators probed competitive implications, a notably distinct conversation emerged—one that did not flow directly from productivity assessments. Directors who agreed on task-level capabilities diverged on what those capabilities meant for competitive positioning, firm viability, and sectoral structure.

Cross-sector heterogeneity in perceptions of sectoral stability. Comparing across sub-sectors, we observed marked variation in consensus regarding sector-level disruption. In regulated professional services—law and financial advisory—directors generally agreed that sectors would remain structurally intact, even as they debated firm-level advantage. The strongest arguments centered on key client requirements: professional accountability, regulatory certification, output accuracy carrying legal weight. In contrast, strategy consulting exhibited less consensus on sectoral stability, with directors articulating different conceptions of core value propositions. PR and communications directors exhibited still greater anxiety, searching actively for defensibility. The heterogeneity in value proposition conceptions appeared to generate heterogeneity in expectations about competitive implications.

Taken together, these observations provided the foundation for our Stage 2 survey instrument. We designed the survey to quantify constructs that had emerged as salient: directors’ assessments of leverage potential, displacement expectations, views on sectoral disruption, and the organizational and sectoral characteristics they emphasized when explaining these assessments.

5. Stage 2: Structured Survey and Analyses

5.1 Methodology

We designed a survey instrument to examine these patterns across a broader population of directors than roundtables could accommodate. We fielded a Qualtrics instrument to the IoD membership in February–March 2024. Five-point Likert items captured directors' assessments across three domains: perceived competitive implications of GenAI on their sector, expectations regarding firm displacement and competitive differentiation, and the importance of various assets to their firm's success—including proprietary data, tacit knowledge, technical expertise, relational capital, and regulatory certification. The survey yielded 221 complete responses (~11% response rate), all from director-level or higher executives. Randomized item order, centered scale anchors, and attention checks minimize common-method bias. Full methodological details appear in Appendix B.

5.2 Survey Findings

Our analysis proceeded in two phases: descriptive examination of response distributions, followed by ordinary least squares (OLS) regressions to identify systematic associations between firm characteristics and directors' competitive expectations. We summarize key findings here.

Table 3: Survey assessments of organizational characteristics and GenAI's expected impact (Stage 2)

	Mgmt Consult	Software	Banking	Insurance	Education	Manufact	Construct	Overall
Panel A: Org Characteristics								
Proprietary Data ^a	2.22 (1.29)	2.23 (1.01)	2.08 (0.90)	2.00 (1.50)	1.80 (1.01)	2.55 (1.21)	2.50 (1.08)	2.30 (1.21)
Tacit Knowledge ^b	4.42 (0.91)	4.31 (0.75)	4.17 (0.58)	4.67 (0.50)	4.40 (0.63)	4.18 (0.87)	4.60 (0.52)	4.35 (0.79)
Relationship Knowledge ^b	4.58 (0.81)	4.38 (0.96)	4.92 (0.29)	4.67 (0.71)	4.53 (0.64)	4.36 (0.81)	4.70 (0.48)	4.45 (0.84)
Technical Knowledge ^b	4.50 (0.77)	4.38 (0.65)	4.42 (0.51)	4.44 (0.73)	4.33 (0.90)	4.64 (0.67)	4.40 (0.97)	4.28 (0.92)
Market Insight ^b	4.28 (0.81)	4.31 (1.03)	4.33 (0.65)	4.11 (1.05)	4.13 (0.74)	4.09 (0.70)	4.60 (0.70)	4.12 (0.91)
Certification ^c	2.94 (1.53)	2.77 (1.48)	1.50 (0.52)	1.44 (0.53)	2.20 (1.21)	2.45 (1.57)	2.00 (1.05)	2.47 (1.42)
Regulation Level ^d	3.14 (1.13)	3.00 (1.29)	4.33 (0.65)	4.67 (0.50)	3.73 (0.96)	3.64 (1.03)	3.80 (0.79)	3.57 (1.04)
Modularity ^e	2.11 (0.89)	1.62 (0.77)	2.33 (0.65)	2.22 (1.09)	1.80 (0.56)	2.45 (0.69)	1.70 (0.48)	1.99 (0.76)
Pattern Recognition ^f	2.72 (1.26)	2.15 (1.07)	2.33 (0.89)	2.22 (0.97)	2.13 (0.99)	2.82 (0.98)	2.10 (0.88)	2.39 (1.15)
Panel B: Competitive Implications								
GenAI's Industry Importance (5yr) ^g	3.89 (1.04)	3.92 (1.26)	3.83 (0.83)	3.89 (1.05)	4.27 (0.70)	3.45 (0.93)	3.60 (1.26)	3.82 (1.02)
Distinctiveness Potential ^h	5.33 (1.15)	5.69 (1.32)	5.42 (1.24)	5.00 (0.87)	5.87 (0.74)	5.09 (1.58)	5.30 (1.25)	5.31 (1.37)
Difficulty: Competitor Imitation ⁱ	2.72 (1.09)	3.00 (1.00)	3.75 (0.97)	2.56 (1.13)	3.73 (0.70)	2.45 (1.21)	3.70 (1.06)	3.00 (1.20)
Difficulty: Supplier Substitution ⁱ	2.72 (1.19)	3.00 (1.00)	3.25 (0.97)	2.00 (0.87)	3.73 (0.59)	2.64 (1.29)	2.60 (1.26)	2.85 (1.15)
Difficulty: Customer Bypass ⁱ	2.61 (1.10)	2.54 (1.27)	3.17 (0.94)	2.11 (1.05)	3.53 (0.92)	2.27 (1.27)	3.00 (1.15)	2.73 (1.23)
Likelihood: Financial Viability ^j	3.33 (1.04)	3.08 (1.32)	3.67 (0.98)	3.11 (1.45)	3.87 (1.13)	2.73 (1.01)	3.00 (1.05)	3.25 (1.13)
Panel C: Enhancement Potential								
Enhance Products/Services ^k	3.47 (1.08)	3.69 (1.38)	3.50 (0.90)	2.56 (1.24)	3.80 (0.94)	2.91 (1.30)	3.40 (1.43)	3.37 (1.21)
Build New Products ^k	3.06 (1.39)	3.54 (1.51)	3.33 (0.98)	2.11 (1.27)	3.73 (1.10)	3.00 (1.18)	2.80 (1.48)	3.06 (1.42)
Enhance Workflows ^k	3.44 (1.05)	3.69 (1.32)	4.25 (0.62)	3.89 (1.17)	4.00 (0.76)	3.36 (1.12)	3.30 (0.95)	3.55 (1.13)
Enhance Productivity ^k	3.17 (1.11)	3.31 (1.49)	3.92 (0.67)	3.44 (1.33)	3.73 (1.03)	3.36 (1.29)	3.30 (1.16)	3.44 (1.13)
Enhance Creativity ^k	3.42 (1.13)	3.62 (1.39)	3.50 (1.00)	2.44 (1.42)	4.20 (0.94)	3.27 (1.01)	3.70 (0.95)	3.38 (1.17)
<i>n</i>	36	13	12	9	15	11	10	221

Notes: Standard deviations in parentheses. All items measured on 1–5 scales unless noted. (a) Proprietary data importance. (b) Knowledge importance (1 = Not important, 5 = Critically important). (c) Certification importance. (d) Regulation level (1 = Low, 5 = High). (e) Modularity (1 = Highly integrated, 3 = Highly modular). (f) Pattern recognition importance. (g) Industry importance of GenAI (5yr). (h) Distinctiveness potential (1–7 scale). (i) Displacement difficulty (1 = Very easy, 5 = Very difficult). (j) Challenge likelihood (1 = Very unlikely, 5 = Very likely). (k) Enhancement potential (1 = Not at all, 5 = A great deal)."

First, directors generally agreed that GenAI will prove competitively consequential—as seen in **Table 3**, 89% anticipated at least moderately significant impact on their industry over the next five years. Yet this consensus masked the heterogeneity that Stage 1 had surfaced qualitatively. When we disaggregated “competitive implications” into expectations about firm displacement versus competitive differentiation, response distributions diverged markedly. Differentiation potential exhibited a unimodal distribution with positive skew—directors generally perceived opportunity. Displacement expectations approached uniformity, with substantial minorities anticipating both greater ease and greater difficulty of displacing incumbents. Within narrowly defined sub-sectors, directors exhibited marked disagreement on competitive implications. In management consulting ($n=36$), standard deviations on disintermediation likelihood ($SD=1.25$) and distinctiveness potential ($SD=1.15$) approached the overall sample variance, indicating that sub-sector membership explains little of the observed heterogeneity. Banking/investment directors showed greater consensus on competitive outcomes (disintermediation $SD=1.08$; distinctiveness challenge $SD=.74$), while tech directors exhibited the highest variance (distinctiveness potential $SD=1.32$).

Second, OLS results reveal that within-sector heterogeneity associates systematically—but differentially—with how directors see their organizations. As seen in **Table 4**, no single feature predicts all dimensions of competitive expectations; rather, different features matter for different dimensions. Pattern-recognition importance comes closest to a consistent predictor: directors rating it as central to their success perceived lower differentiation potential ($b = -.22, p < .05$), lower disintermediation likelihood ($b = -.20, p < .05$), and lower financial viability risk ($b = -.17, p < .05$). Modularity uniquely predicted industry importance ($b = -.24, p < .05$), while regulation uniquely predicted differentiation potential—positively ($b = .32, p < .05$), suggesting directors in regulated industries perceive GenAI as an opportunity within protected boundaries. These associations hold with sector fixed effects.

Third, cross-sectoral comparisons corroborated Stage 1’s observation that regulated sectors exhibit greater consensus. Directors in highly regulated industries showed markedly lower variance on distinctiveness potential ($SD = 1.20$ vs. 1.87) and GenAI’s expected importance for their industries ($SD = .94$ vs 1.27). Banking and insurance directors converged on competitive expectations; management consulting and software directors exhibited wider dispersion, with standard deviations on distinctiveness potential approximately 1.4 times larger. Notably, we do not find evidence that these reflect differences in productivity evaluations—assessments of GenAI’s potential for productivity enhancement exhibited similar variance across regulation levels.

Table 4: OLS Predictors of Competitive Expectations Regarding GenAI (Stage 2)

	Industry Importance (a)			Distinctiveness Potential (b)			Disintermediation Likelihood (c)			Financial Viability Risk (d)		
	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)	(1)	(2)	(3)
Proprietary Data	-0.230 (0.000)	-0.108 (0.131)	-0.051 (0.497)	-0.259 (0.001)	-0.120 (0.222)	-0.064 (0.558)	-0.290 (0.000)	-0.133 (0.125)	-0.071 (0.455)	-0.271 (0.000)	-0.192 (0.031)	-0.151 (0.126)
Tacit Knowledge	0.266 (0.015)	0.028 (0.795)	0.012 (0.917)	0.239 (0.041)	0.272 (0.114)	0.264 (0.150)	0.029 (0.796)	-0.196 (0.107)	-0.227 (0.076)	0.048 (0.655)	-0.096 (0.494)	-0.105 (0.480)
Technical Knowledge	0.303 (0.001)	0.114 (0.241)	0.093 (0.343)	0.279 (0.006)	0.092 (0.439)	0.083 (0.513)	0.151 (0.098)	-0.057 (0.554)	-0.051 (0.653)	0.126 (0.198)	-0.028 (0.790)	-0.028 (0.807)
Pattern Recognition	-0.201 (0.003)	-0.060 (0.411)	-0.055 (0.466)	-0.312 (0.000)	-0.220 (0.009)	-0.216 (0.016)	-0.340 (0.000)	-0.206 (0.010)	-0.198 (0.021)	-0.243 (0.001)	-0.175 (0.027)	-0.169 (0.040)
Modularity	-0.387 (0.000)	-0.250 (0.022)	-0.244 (0.039)	-0.276 (0.037)	-0.111 (0.408)	-0.124 (0.412)	-0.299 (0.007)	-0.155 (0.189)	-0.151 (0.252)	-0.160 (0.142)	-0.080 (0.458)	-0.068 (0.562)
Regulation	0.131 (0.083)	0.117 (0.178)	0.135 (0.153)	0.216 (0.052)	0.262 (0.051)	0.319 (0.037)	0.184 (0.024)	0.101 (0.282)	0.119 (0.271)	0.091 (0.247)	-0.003 (0.975)	-0.020 (0.853)
Relationship Knowledge	0.217 (0.045)	0.032 (0.739)	-0.006 (0.958)	-0.058 (0.629)	-0.345 (0.044)	-0.298 (0.123)	0.120 (0.306)	0.079 (0.575)	0.096 (0.486)	0.180 (0.066)	0.166 (0.208)	0.096 (0.478)
Market Insight	0.339 (0.000)	0.136 (0.166)	0.195 (0.077)	0.330 (0.001)	0.167 (0.183)	0.185 (0.201)	0.306 (0.001)	0.165 (0.199)	0.238 (0.098)	0.187 (0.035)	0.011 (0.918)	0.029 (0.819)
Certification	-0.080 (0.121)	0.009 (0.873)	-0.001 (0.991)	-0.091 (0.169)	0.050 (0.577)	0.049 (0.655)	-0.200 (0.000)	-0.125 (0.067)	-0.147 (0.050)	-0.129 (0.018)	-0.030 (0.681)	-0.058 (0.474)
Referrals	0.208 (0.001)	0.017 (0.811)	0.018 (0.813)	0.104 (0.212)	-0.046 (0.698)	-0.104 (0.472)	0.166 (0.029)	0.029 (0.775)	0.018 (0.871)	0.080 (0.308)	-0.032 (0.740)	-0.028 (0.791)
Firm Size (Turnover)	-0.065 (0.142)	-0.080 (0.106)	-0.045 (0.385)	-0.103 (0.079)	-0.115 (0.059)	-0.064 (0.391)	-0.087 (0.134)	-0.138 (0.017)	-0.128 (0.037)	0.056 (0.281)	0.020 (0.712)	0.063 (0.307)
<i>Sector FE</i>	No	No	Yes	No	No	Yes	No	No	Yes	No	No	Yes
<i>R²</i>		0.206	0.331		0.187	0.258		0.216	0.277		0.130	0.221
<i>Observations</i>		221	221		221	221		221	221		221	221

Notes: OLS coefficients with exact p-values in parentheses; robust (HC3) standard errors used throughout. Columns (1) report bivariate associations. Columns (2) include all predictor and control variables. Columns (3) add sector fixed effects. Dependent variables: (a) “How important do you perceive the usage of GenAI to be within your industry within the next 5 years?” (1–5 scale). (b) “Assess the potential for GenAI to unlock new opportunities for distinctiveness within your market” (1–7 scale). (c) “Rate the likelihood to which you expect GenAI to challenge disintermediation within your organization over the next 5 years” (1–5 scale). (d) “Rate the likelihood to which you expect GenAI to challenge financial viability within your organization over the next 5 years” (1–5 scale).

5.3 Emergent Patterns and Transition to Stage 3

Stage 2 confirms quantitatively that directors agree on GenAI's competitive salience while diverging on what that salience entails—and that this divergence associates systematically with how they regard their organization's assets. Yet, survey data document that these associations exist without illuminating why. When directors rate tacit knowledge as central to their firm's success, what reasoning connects that assessment to their competitive expectations? Moreover, the finding that technical expertise associates with both opportunity and threat—while relational capital associates primarily with insulation—suggests more complex interpretive dynamics than the survey can capture. What is it about how directors understand their firm's work that leads some to perceive GenAI as reinforcing their position, while others anticipate vulnerability?

6. Stage 3: Insights from Follow-Up Roundtables

Stage 2 documented that directors' competitive expectations associate with how they see their firms' assets. Stage 3 was designed to examine why. To do this, we inverted Stage 1's design logic: rather than grouping directors by sector to surface disagreement, we used Stage 2 survey responses to construct roundtables matching directors who shared similar response patterns. This design served two purposes. First, it enabled us to probe the reasoning underlying different expectations by enabling deeper discourse amongst participants with similar expectations. Second, by distributing directors from the same sub-sector across different roundtables rather than clustering them together, we could observe whether they independently articulated similar reasoning or diverged in their interpretations even when not directly confronting one another.

We conducted nine roundtables with 163 directors between March and May 2024. Seven paired roundtables created matched comparisons along key dimensions from Stage 2; two additional roundtables with private equity executives provided an outside-in perspective. Methodological details and coding procedures appear in Appendix C; here we present findings organized around the three dimensions of competitive implications that emerged from our investigation.

A preliminary observation warrants emphasis. Directors did not attend to only one dimension of competitive implications to the exclusion of others. Rather, they appeared to conceptualize these dimensions as distinct phenomena requiring separate assessment. The same director who expressed confidence about leveraging GenAI for differentiation might simultaneously voice concern about displacement threats to particular competitors, or about the broader trajectory of the industry. Accordingly, we present findings on each dimension in turn, recognizing that individual directors' assessments often spanned all three. See **Table 5** for a detailed breakdown.

Table 5: Director categorization of GenAI's competitive implications (Stage 3)

Dimension	Level of Analysis	Constructs	Representative Quotes	Corresponding Concerns from Directors
Threat of Substitution	Firm Features	Technical Knowledge	A1a (engineering consultancy): “I think in terms of our industry, our industry is all about the application of technical knowledge [...] I think the big fear [is to] render us useless as such and there is a concern about that [...] [the] danger is getting overtaken [...]”	→ Evaluating whether there are areas of technical knowledge that the organization uniquely possesses → Considering reduction of junior workforce and concentrating on hiring staff with higher expertise
		Relational Capital	A2a (restructuring): “Operational efficiency, over time as people are learning, will lead to massive initial transformations because costs will come down significantly and people will be able to use their time on more human-type relationship areas of focus.”	→ Questioning how relational capital is developed when client-facing exposure diminishes → Examining whether increased in-house capabilities erode the touchpoints through which trust-based relationships are traditionally built
	Sectoral Features	Regulatory Intensity	A2a (banking): “We [in financial services] will test everything because we are regulated. [...] So, one, of course, is the AI regulation itself. How would FSFCA, who regulates us, see it? I think there's not one regulation where we can talk and find things.”	→ Monitoring evolving AI-specific regulation and the shifting boundaries of permissible automation → Recognizing increased organizational impetus to hire specialized expertise for output verification, given that accuracy carries regulatory and reputational stakes
		Sectoral Norms	A1b (financial services): “I think there is [...] a culture around fear of the new technology and what it means. [...] But I do believe that those who use it, who actually take the risk [...] might have an advantage against your competitors.”	→ Assessing whether institutionalized practices are genuinely stable or susceptible to disruption → Heightened vigilance toward startups and leaner entrants unburdened by legacy commitments
Leverage GenAI for Differentiation	Firm Features	Proprietary Data	B2 (law firm): “You can't just put that into ChatGPT and get an answer to draft a French energy nuclear contract. It's impossible because the system isn't trained on that specific data, which likely exists only in a few cloud-based files.”	→ Shifting attention toward knowledge management and the curation of internal “archives” → Focusing on the labor-intensive tasks required to make historical data “useful” for new models
		Tacit Knowledge	A2b (construction): “In our sector, tacit knowledge is the intricacies behind building a wall and then plastering it. [...] Our expectation is that we have the people with the best understanding of these tacit features and hence will be able to use GenAI better than others.”	→ Identifying “edge-case” tasks where human judgment and physical experience remain irreplaceable → Questioning how tacit knowledge can be cultivated when the apprenticeship processes through which it was traditionally developed are increasingly automated
		Flexibility to Experiment	A3a (manufacturing): “The ability to experiment freely with new technologies like AI isn't something you can just decide to have. [...] It's built up over years, maybe decades, of fostering a culture that embraces risk and learning from failure. [...]”	→ Benchmarking against organizational cultures perceived as more adaptive to technological change → Attending to startups and new entrants unencumbered by legacy systems and reputational liabilities, who can iterate more freely
		Data Cleanliness	A1a (technology firm): “We have all this proprietary data that's incredibly valuable, but accessing it is a nightmare. The effort to merge and clean the data is enormous, and we're not even sure if we can do it. It's a major obstacle in utilizing our data to its potential.”	→ Focusing on the “plumbing” of organizational systems and the effort required to merge legacy data → Identifying the specific bottlenecks in data access that hinder task-level automation
Industry Disruption	Value Proposition	Modularity	A4 (financial services firm): “We have tried to strive for a more modular structure. We're acutely aware that GenAI could potentially replace entire functional units. This modularity makes us more adept, more agile, but also more vulnerable to plain imitation.”	→ Evaluating which functional units are truly independent versus those that are tightly integrated → Looking into their service offering and whether they can bundle things differently
		Pattern Recognition	B1 (consulting): “If you strip away the branding, most of what we do at this level is just high-level pattern recognition—looking at a client's mess and matching it to something we've seen a dozen times before [...] and [...] that's exactly what GenAI is built to do.”	→ Shifting focus toward high-level synthesis rather than the routine collection of patterns → Re-evaluating the “value-add” of junior analysts and the downsizing of lower layers of the hierarchy

6.1 Leveraging GenAI for Competitive Differentiation

Directors who perceived opportunity for differentiation emphasized complementary assets they believed would become more valuable alongside GenAI capabilities. Three categories of such assets recurred across our roundtables: proprietary data, tacit knowledge, and an organizational culture that supported experimentation.

Proprietary data was the most frequently cited source of differentiation potential. Directors described recognizing latent value in hitherto neglected information assets. A law firm partner captured this view: “You can’t just put a French energy nuclear contract into ChatGPT. That data exists only in a few cloud-based files—it’s incredibly valuable, high-end, and rare. And we have billions of such data points.” A consulting firm managing partner described a similar realization: “We’ve been collecting data from thousands of client engagements over decades, but we’ve barely scratched the surface of what insights we might unlock.” These directors perceived GenAI not as a substitute for their accumulated knowledge but as a tool that could finally extract value from assets that had been difficult to leverage through conventional means.

Tacit knowledge represented a second category. Directors emphasized expertise that could guide effective deployment and interpretation of GenAI outputs—knowing which questions to ask, how to evaluate responses, and where human judgment remained essential. An accountancy director framed this as “the ability to distinguish between Type I and Type II errors—knowing when the machine is confidently wrong versus usefully right.” This knowledge, directors argued, could not be easily codified or replicated. Two directors in separate roundtables independently argued that GenAI, by automating the very tasks through which tacit knowledge had traditionally been developed, might paradoxically make such knowledge more scarce and therefore more valuable as a source of competitive differentiation. These directors spoke about ongoing efforts to restructure their organizations in ways that would preserve the development and transmission of tacit knowledge even as routine tasks became automated.

Flexibility to experiment constituted the third category. Some directors pointed to cultures that permitted rapid experimentation as a source of advantage, arguing that the ability to iterate quickly on GenAI applications would prove more valuable than any specific technical capability. A financial services director noted: “The firms that will win aren’t necessarily the ones with the best models—they’re the ones willing to try things, get them wrong, learn, and try again.” Several directors acknowledged that their own firms lacked such cultures. Unprompted, these directors named specific competitors within their sub-sector whom they perceived as better positioned to experiment. A strategy consulting director observed: “We’re too risk-averse, too concerned with maintaining our brand for precision. I can think of two or three firms in our space—you’d know the names—who have the right DNA for this. They’ll figure out how to use it before we do.”

6.2 Threat of Firm Displacement

Assessments of displacement threats associated with how directors saw the nature of expertise—their own and competitors’. The crucial distinction was the codifiability of knowledge.

Directors referencing firm propositions grounded in **codifiable technical knowledge**—expertise acquired through formal training and amenable to explicit articulation—expressed concern that GenAI could substitute this expertise and erode competitive position. As one engineering consultancy director put it: “Our industry is all about the application of technical knowledge [...] the big fear is that larger companies may be able to leverage AI and render us useless.” A strategy consulting director was blunter: “If your business still relies on technical knowledge, then you’re toast.” But these concerns were not always self-referential. Several directors described displacement threats as applying primarily to competitors whose value propositions they regarded as more technically oriented than their own. A consulting director observed: “The boutiques that compete with us on deep functional expertise—supply chain optimization, pricing analytics—they’re the ones exposed. What we do is different; it’s about navigating organizational politics, managing stakeholder alignment. That’s not getting automated anytime soon.” Another director named specific firms: “Firms like [X] and [Y] have built their reputation on analytical horsepower. That’s exactly what the models are getting good at. I wouldn’t want to be them right now.”

Directors who emphasized **relational capital** or **regulatory complexity** perceived greater insulation—again, sometimes for themselves, sometimes for competitors. Their reasoning was not that GenAI lacked relevant capabilities, but that value propositions resided in dimensions that GenAI could not replicate. A restructuring executive captured this view: “A lot of your work is around knowledge, data, content [...] so operational efficiency will lead to massive initial transformations because costs will come down significantly and people will be able to use their time on more human-type relationship areas of focus.” The implication was that firms whose value derived primarily from knowledge and content production would face displacement, while those whose value derived from relationship management would be repositioned rather than threatened.

Directors in regulated sectors pointed to different sources of insulation that operated at the sectoral level rather than the firm level. The distinguishing features were not simply the presence of regulation but specific characteristics of what clients required. Directors emphasized contexts in which clients needed not merely high-quality work but verified, certified, or professionally warranted outputs—contexts in which being 99 percent correct was categorically different from being 100 percent correct. An audit partner explained: “Our clients don’t pay us to produce financial statements. They pay us to sign off on them, to put our name and our liability on the line. That’s what they’re buying—the certification, not the analysis.” A legal services director made a similar observation: “In litigation, you can’t be approximately right. You need case citations that exist, precedents that hold up, arguments that won’t embarrass you in court. The stakes of error are too high.” These sectoral features associated with directors expressing less concern about displacement for firms operating within such regulatory structures.⁴

⁴ Roundtable dynamics may have contributed to directors’ willingness to discuss competitors’ vulnerabilities alongside their own positions. In group settings, directors appeared more comfortable articulating threats in terms of “firms like X” or “the boutiques that do Y” than in terms of their own organizations alone. Whether this reflects genuine comparative assessment or a more performative distancing from acknowledged vulnerabilities, we cannot determine from the

6.3 Broader Industry Disruption

Across all nine roundtables, a common pattern distinguished directors who expressed concern about industry-level disruption from those whose concerns remained at the firm level. Two features consistently associated with the locus of concern: whether directors regarded their work as **modular**—decomposable into separable components that could be performed independently—and whether they regarded it as **reducible to pattern recognition**—the identification of regularities in data that machines might replicate.

What proved revealing was not simply whether directors perceived these features as present or absent, but where they located them. Directors who perceived modularity or pattern-recognition reliance as distinctive to their own firm expressed concern about firm-level displacement, while those who perceived these same features as characteristic of their sector as a whole expressed concern about broader industry disruption.

The firm-level framing surfaced most clearly among directors who had made deliberate architectural choices they now perceived as sources of exposure. A director of a mid-sized consultancy illustrated this pattern: “Our competitors have much more integrated offerings—clients come to them for the relationship, the full package. We’ve deliberately modularized, made ourselves efficient by unbundling. That efficiency has been our advantage, but now I worry it makes us specifically easier to replicate. The big integrated players will be fine. We’re exposed.” The concern here was positional: GenAI threatened to erode this particular firm’s standing relative to differently structured competitors. Other directors in similar positions described analogous reasoning—they had optimized for efficiency through decomposition, and now worried that the very modularity that enabled efficiency also enabled replication. Their attention remained on firm-level competitive dynamics, on whether their organization would gain or lose relative position.

The sector-level framing surfaced among directors who perceived modularity as an industry-wide structural feature rather than a firm-specific choice. A professional services director captured this view: “It’s not that we’re particularly vulnerable—everyone in this industry bills by the hour for work that can increasingly be broken into discrete deliverables. [...] So when I think about AI, I’m not worried about losing clients to competitors. I’m worried that the whole model—the way professional services have worked for decades—gets unbundled. It’s not a firm problem. It’s a sector problem.” Directors who framed modularity at the sector level expressed concern about fundamental restructuring of how value was created and captured in their industries—concern that transcended any particular firm’s competitive position.

The pattern-recognition dimension operated similarly. Directors who described their own firm’s work as particularly reliant on identifying patterns in data worried about firm-level displacement: “We’ve built our reputation on finding the signals in messy data sets. That’s literally what machine learning does. Other firms in our space do more bespoke advisory work—they’ll adapt. We’re the ones who need to reinvent ourselves.” Directors who described pattern recognition as characteristic of their sector worried about industry-level transformation: “Financial analysis

roundtable data. Regardless, directors’ displacement assessments frequently encompassed the competitive landscape rather than remaining narrowly self-referential.

has always been about pattern recognition—reading the signals, spotting the trends. Every firm in this industry does some version of that. The question isn't which firm gets disrupted; it's whether the whole industry gets restructured when machines can do the pattern-finding faster and cheaper than any of us."

Furthermore, within sub-sectors, directors often diverged on where they located these features. Even among directors matched on firm size and customer offering, some described their own organizations as having "processes so interconnected that it's hard to see how GenAI could simply slot in," while perceiving competitors as more modular. Others described their own firms as deliberately unbundled for efficiency, while viewing the sector's integrated incumbents as insulated by their complexity. This variation associated systematically with which dimension of competitive implications directors emphasized. Those who located modularity and pattern-recognition reliance at the firm level attended to displacement; those who located these features at the sector level attended to industry disruption⁵. That said, as we explain in the Appendix, the relationship between structural perception and competitive expectation was not uniform.

⁵ To be clear, informants did not use such theoretical terms explicitly. One director described "connecting dots across different data sources, finding the anomalies" rather than invoking "pattern recognition"; another explained that "you can't just hand off pieces of this to someone else—it all has to fit together in your head" rather than discussing "modularity." This linguistic distance between practitioner accounts and analytical constructs suggests these cognitive and structural dimensions operate tacitly, embedded in professional practice rather than consciously articulated.

7. Stage 4: Survey, Interviews, and Roundtable, 18 Months On

The patterns documented in Stages 1–3 emerged from data collected 15–18 months after ChatGPT's initial launch, when uncertainty remained high and few directors had made public commitments constraining candid assessment. GenAI capabilities expanded substantially through 2024 and into 2025—new model releases, expanding enterprise adoption, accumulating evidence on productivity effects (e.g., Handa *et al.*, 2025). A natural question is whether the patterns we documented reflect durable cognitive frames or transient reactions to a novel technology. Stage 4 was designed to examine this question through two methods: a newly designed survey instrument administered to 41 directors—27 of whom had participated in Stage 2, enabling cross-wave comparison over a 14-month interval—and semi-structured individual interviews with 22 directors, a method deliberately chosen to avoid the peer-group dynamics that might have reinforced earlier positions in roundtable settings. Survey and interview protocols appear in Appendix D.

A newly designed survey administered to 41 IoD directors probed the enhancement–displacement distinction directly. Directors rated the likelihood that GenAI could enhance their core offerings substantially higher than the likelihood it could replace them ($M = 3.78$ vs. 2.44 ; paired $t(40) = 6.35$, $p < .001$, $d = 1.00$); two-thirds rated enhancement as more likely, while only one director rated displacement higher. When asked to imagine that GenAI could automate 50 percent of professional tasks, directors still rated expected competitive erosion below the scale midpoint ($M = 2.59$, $p = .013$), with 83 percent anticipating moderate erosion or less. Value proposition assessments revealed a complementary pattern: directors rated relationship management as most important to current value creation ($M = 4.20$) yet expected it to receive the lowest GenAI usage over five years ($M = 3.07$; gap = 1.12 , $p < .001$), while process efficiency—rated moderately important—received the highest expected usage.

The predictors of these expectations exhibited a consistent asymmetry. Directors who spent more time with GenAI perceived greater enhancement potential ($r = .60$, $p < .001$) but not greater displacement risk ($r = .29$, $p = .064$); organizations that had created GenAI-specific roles—29 percent of the sample—reported dramatically higher enhancement expectations ($d = 1.43$, $p < .001$) while their displacement expectations did not differ. Direct experience selectively amplified opportunity beliefs without revising threat beliefs. Among 27 directors who had also completed the Stage 2 survey 14 months earlier, this selectivity extended across waves: Stage 2 ratings of technical knowledge importance predicted Stage 4 enhancement expectations ($r = .54$, $p = .004$), and Stage 2 ratings of relationship knowledge importance were associated with lower Stage 4 displacement expectations ($r = -.34$, $p = .079$).

Individual interviews imparted texture to these patterns. Directors uniformly acknowledged that GenAI capabilities had advanced substantially. A managing director at a strategy consultancy was representative: “Eighteen months ago, it was impressive for what it was—a sophisticated autocomplete. Now it can genuinely reason through complex problems. The trajectory has been steeper than I expected.” Several directors also acknowledged concrete workforce implications—an accounting director noted that successive waves of automation had reduced his team from six

juniors to one over three decades, with GenAI the latest iteration: "All the automation has happened at the lower end. It's the thinking part, the judgment, that's popped out at the top." Yet, when probed on whether these improvements had altered their competitive expectations, directors overwhelmingly described stability rather than revision. The same managing director continued: "But does that change whether clients need us? Not really. If anything, I'm more confident than I was. We've had 18 months to see how clients actually use this [...], they use it for the analytical grunt work, just as we do, but they still need someone to tell them what it means." Directors who had expressed concern maintained their concern; those who had perceived insulation continued to perceive it. A corporate law partner captured the reasoning: "The technology is better, obviously. But the fundamental question hasn't changed—can it be liable? Can it carry professional indemnity? Can clients sue it when it's wrong? Those questions aren't about capability."

How did directors reconcile acknowledged capability expansion with stable competitive expectations? Three patterns recurred. Directors who had emphasized proprietary data or tacit knowledge as sources of differentiation saw capability improvements as amplifying existing advantages—a director at a marketing agency described consolidating forty years of project documentation into an AI-powered retrieval system, noting "you still can't do it without an engineer, but the supporting information appears at the touch of a button." Better models meant better leverage on accumulated assets that competitors could not replicate by subscribing to the same tools. Directors who had emphasized relational or implementation dimensions described capability expansion as irrelevant to what they perceived as the dimensions that mattered; as one restructuring executive put it, "Whether the AI can analyze a balance sheet in ten seconds or one second doesn't change whether it can sit across from a management team and convince them to take difficult decisions." Directors who had located concerns at the sector level maintained those concerns despite firm-level adaptations; they continued to question whether their industry's fundamental architecture—billing models, value chain positions, client relationships—would survive, regardless of how individual firms responded. Across all three patterns, directors updated their views on what GenAI could do without correspondingly updating their views on what it meant competitively—suggesting that the interpretive frames through which they filtered GenAI's implications proved more stable than their assessments of the technology itself.

8. Stage 5: Interviews, Two Years On

Stage 5 interviews were conducted one-to-one in January–March 2026 with 41 directors who had participated in at least one prior stage. Stage 5 examined whether the temporal stability documented in Stage 4 persisted as GenAI moved from early experimentation toward routinized organizational deployment. The individual interview format remained consistent with Stage 4's methodological shift away from group settings. See Appendix E for details.

The central finding was continuity: directors' competitive expectations remained consistent with their earlier assessments, reinforcing the decoupling between capability beliefs and competitive frames documented in Stage 4. Directors again acknowledged substantial advances in GenAI's technical performance—several pointed to improvements in agentic reasoning, multi-step task execution, and domain-specific fine-tuning as developments they had not anticipated—yet these acknowledged advances did not translate into revised competitive expectations. Directors who had emphasized proprietary data, relational capital, or regulatory buffers continued to invoke these same features; those who had expressed displacement or disruption concerns maintained them.

What did shift, however, was the focus of directors' attention. While earlier stages had been dominated by forward-looking assessments of whether GenAI would prove competitively consequential, Stage 5 conversations increasingly turned toward diagnosing why competitive impact had not materialized as rapidly as anticipated. Several directors expressed a degree of surprise—and, in some cases, frank disappointment—at the gap between GenAI's demonstrated task-level capabilities and observable competitive consequences. A technology consulting director captured this sentiment: "I'm a little disappointed, both in myself and the whole setup—because I would have expected something by now." Another consulting director observed that despite widespread adoption, "I've not seen anybody in a large organization make a case for restructuring and rationalizing the workforce based on the availability of LLMs." This attentional reorientation—from anticipating impact to interrogating its absence—marked a notable shift in the tenor of directors' engagement with GenAI's competitive implications, even as their underlying expectations for the technology's eventual competitive impact remained stable.

In seeking to explain this gap, directors within the same sub-sectors began to articulate different accounts of where and why competitive effects were materializing—or failing to materialize. Here, modularity and the structural decomposability of value propositions emerged as a considerably more salient and *explicit* dimension than in earlier stages. Directors increasingly distinguished between service areas that were relatively self-contained and separable from the broader offering—and therefore more susceptible to client insourcing or substitution—and those that were deeply integrated with other components of the value proposition. As one management consultant put it, "What used to be the consultant's domain, now the clients can do themselves"—referring specifically to data-gathering and preliminary analysis. Several directors described observing clients taking discrete analytical functions in-house, particularly where those functions operated independently of relationship-intensive or judgment-dependent aspects of the engagement. A private equity director made an analogous observation about legal services: tasks governed by

narrow, well-defined rules—what he described as areas where “the rules of engagement are very small”—were those where GenAI deployment was advancing most rapidly.

More broadly, directors acknowledged that the problem-solving and pattern-recognition components of their offerings were eroding as sources of differentiation, with one director characterizing GenAI’s effect on analytical output as producing “vanilla”—competent but undifferentiated work that clients could increasingly generate without professional intermediation. The aspects of the business that directors scrutinized most closely were those that were truly independent from other service components: modular, analytically oriented deliverables that could be separated from the bundled value proposition without disrupting the whole. Directors whose offerings were more tightly integrated—where analytical work was embedded within relationship management, implementation oversight, or regulatory compliance—expressed comparatively less concern about these dynamics, though they were not immune to broader sectoral anxieties.

Taken together, Stage 5 extends the temporal window of our documentation and reveals a subtle but consequential shift: while the ordering of competitive expectations remained stable, the interpretive work directors performed around those expectations evolved. Directors were actively theorizing about the structural features of their value propositions that explained why certain competitive effects were emerging while others were not.

9. Discussion

This study documents how corporate directors interpret GenAI's competitive implications as they unfold. Across survey responses, interviews, and successive roundtables, directors consistently treated GenAI's competitive implications along three dimensions: (1) the potential to leverage GenAI for *competitive differentiation*, (2) the threat of *firm displacement*, and (3) broader *industry disruption*. Importantly, directors did not simply offer “more optimistic” or “more pessimistic” views: expectations along one dimension did not necessarily correspond with expectations along the others. The result is a patterned heterogeneity that reflects how directors construe their firms' and sectors' value propositions and complementary assets, not only the tasks GenAI can perform.

What makes this compartmentalization consequential is the logic directors used to justify it. Even within narrowly defined sub-sectors, directors invoked different organizational features when interpreting GenAI's implications. Some emphasized proprietary data, accumulated expertise, or tacit knowledge as sources of differentiation—arguing that stronger models amplify advantages difficult to replicate. Others cited sector-level conditions—especially regulation, liability regimes, certification demands, and institutional conservatism—to reason about insulation or disruption. Directors within the same sub-sector often disagreed sharply about whether the value proposition is best described as modular versus integrated, and as reducible to pattern recognition versus dependent on judgment, responsibility, and relationship-intensive delivery. That disagreement mattered because it shaped where directors located vulnerability and opportunity—at the firm level (inviting displacement logic) or at the sector level (inviting disruption logic).

A second pattern that stands out is the decoupling of productivity assessments from competitive expectations. Directors readily acknowledged GenAI's expanding capabilities and, in many cases, upgraded their assessments of what the technology could do. Yet their competitive expectations—how GenAI would affect competitive advantage and industry evolution—remained strikingly stable. The interpretive frames through which directors filtered GenAI's implications appeared more durable than their views of the technology itself. This has a bearing on strategy research, suggesting that competitive response may not be tightly coupled to technological capability alone; instead, it may depend at least as much on how executives understand their firms' value proposition and complementary assets when they interpret that capability.

9.1 Theoretical Implications

While our contribution is primarily documentary, the patterns we observe speak to several underexplored questions in strategy. First, our evidence reinforces—and extends—work on the durability of strategic frames in the face of novel technologies (Kaplan and Tripsas, 2008; Tripsas and Gavetti, 2000). What is distinctive in our setting is the unevenness of that durability. Directors often updated their views about task-level performance and productivity potential, but their competitive expectations remained comparatively “sticky.” This suggests that durability may not be a single property of “interpretation” in general, but may differ depending on what is being interpreted: task-level implications versus competitive positioning and industry structure.

Understanding why some interpretive elements are quickly revised while others persist may help explain variation in how organizations adapt during technological transition.

Second, we observe an asymmetry in how directors reason about opportunity versus threat, consistent with classic work on issue interpretation (Dutton and Jackson, 1987; Jackson and Dutton, 1988; Staw, Sandelands, and Dutton, 1981). When discussing differentiation potential, directors tended to look inward, emphasizing firm-specific assets and the possibility of leveraging GenAI to strengthen what they already do. When discussing displacement threat, directors were more outward-looking, emphasizing competitor moves, entrants, and shifts in client behavior. This inward/outward asymmetry is notable because both assessments concern the same technology and competitive landscape. It suggests that “threat” and “opportunity” may not only shape the valence of interpretation but also the attention allocation that underlies competitive expectations—where decision-makers search for causal mechanisms and where they anticipate action.

Third, our findings complicate the increasingly influential view of GenAI’s impact as a function of task exposure (Brynjolfsson and Mitchell, 2017; Dell’Acqua et al., 2023; Eloundou et al., 2024; Felten, Raj, and Seamans, 2021). Task-based approaches have been valuable for mapping where GenAI capabilities intersect with occupational profiles. But our documentation shows that directors who broadly agree on what GenAI can do at the task level can still arrive at sharply different judgments about competitive implications. The gap appears to hinge on how directors translate tasks into a theory of the firm’s value proposition and complements—what actually drives differentiation, what is contestable, and what is protected. That translation is not automatic. In this respect, our three-dimensional framework is useful not because it “predicts” outcomes, but because it makes visible the distinct ‘pipes’ and ‘prisms’ (Wu *et al.*, 2014) through which executives are evaluating GenAI’s potential impact and the organizational attributes they treat as relevant to each.

Fourth, our evidence connects directors’ perceived “buffers” to classic ideas about complementary assets and appropriability (Teece, 1986; Tripsas, 1997). Directors who perceived insulation did not typically claim that their work was immune because tasks were complex. Instead, they pointed to complements and constraints: proprietary data access; institutionalized trust; regulatory requirements; liability frameworks; certification and warranting; and relationship-intensive delivery that embeds accountability. Whether these perceived buffers prove durable is an empirical question that our data cannot resolve; regulatory regimes can shift, and new institutional arrangements can emerge. But the fact that directors systematically invoked these features when reasoning about competitive implications suggests that institutional and complement-based analysis may be central to understanding where GenAI’s competitive effects are felt most acutely, and where they are moderated—at least as perceived by senior decision-makers.

Finally, the prominence of modularity in directors’ accounts highlights a tension with longstanding organization and innovation theory. Directors who perceived their offerings as modular—decomposable into discrete tasks—expressed greater displacement concern, while those who perceived their offerings as integrated expressed greater insulation. Yet modularity is also widely

associated with adaptive benefits, including facilitating integration of new technologies (Baldwin and Clark, 2000; Ethiraj, Levinthal, and Roy, 2008), and has been argued to help firms incorporate GenAI (Baldwin, 2024). The result is a managerial dilemma visible in directors' interpretations: designing for easier GenAI integration may be experienced as increasing contestability, even if integration ultimately strengthens the firm. Our data do not adjudicate whether this perceived tradeoff is real, but they do indicate that the perception of such a tradeoff may influence organizational design choices and competitive posture during early diffusion.

9.2 Implications for Practice and Policy

Beyond theory, our findings point to a practical challenge for leaders attempting to craft GenAI strategy in knowledge-intensive settings: the most consequential disagreements may not be about “whether GenAI works,” but about what the organization is, where value is created, and which complements protect or expose that value proposition. The same firm can be described as a bundle of modular tasks (suggesting displacement risk) or as a relationship- and responsibility-intensive system (suggesting insulation and selective augmentation). That framing is not merely descriptive—it is strategic, because it shapes which threats are taken seriously, which opportunities are pursued, and which stakeholders are treated as pivotal.

For policymakers, directors' accounts suggest that regulation shapes not only the pace of GenAI adoption but also who anticipates competitive impact and how. Directors in sectors with stringent liability frameworks, certification requirements, and institutionalized conservatism frequently described these as buffers—contexts requiring verified, certified, or professionally warranted outputs. This underscores a dilemma: regulatory frameworks providing stability during technological transition may inadvertently encourage complacency among protected incumbents, while those enabling rapid disruption may generate instability that overwhelms adaptive capacity. Our data do not resolve this tradeoff, but suggest its contours vary substantially across and within sectors, depending on how actors interpret their complements and constraints.

9.3 Generalizability, Limitations, and Future Research

Our evidence captures director perceptions in a particular context: (mostly) UK corporate directors in knowledge-intensive sectors during GenAI's early diffusion. While we believe the patterns we document are likely to be informative for similar settings (and our out-of-sample roundtables help on this front), we caution against extrapolating to dissimilar institutional environments, or to later stages of maturation when realized competitive consequences become more legible.

A central limitation is that we document how directors see their organizations and environments—not how those environments “objectively” are⁶. We cannot determine whether directors who

⁶ We report two supplementary external-validation exercises in Appendices F-G. First, we link Stage 2 survey responses to a text-based measure of firm differentiation constructed from archived corporate websites (Wayback Machine; Appendix F). Second, we link early-2024 survey expectations from Stage 2 to independently recorded UK firm outcomes (FAME/Bureau

perceive insulation are accurately reading structural features of their competitive contexts or are engaging in motivated reasoning that emphasizes reassuring organizational attributes. This limitation is inherent to perceptual data. At the same time, it points to a valuable agenda for future research: studies that track whether directors' expectations prove accurate—and whether firms acting on different expectations achieve different outcomes—would provide an important complement to our documentation. In that sense, our contribution is best understood as generative rather than conclusive: we map the variables, distinctions, and interpretive linkages that appear salient to senior decision-makers as they reason about GenAI, offering a foundation that future work can refine and extend as evidence accumulates.

The stability we document over time also invites further inquiry. Our window suggests that competitive frames can remain durable even as perceived capability expands, but this durability may not persist indefinitely. As concrete evidence accumulates, directors' frames may eventually shift—or may resist revision, filtering new information through existing interpretive schemas. Longitudinal research across multiple institutional contexts could clarify when interpretive revision occurs, how it spreads within elite networks, and whether certain complements (e.g., proprietary data, regulation, relational capital) become more or less salient as GenAI's deployment moves from experimentation to routinized practice.

More broadly, we document executive reasoning about a technology whose strategic implications remain unsettled. The patterns we identify—distinct dimensions of competitive assessment, within-sub-sector heterogeneity, and the decoupling of productivity judgments from competitive expectations—suggest that executive interpretation of technological change is more textured than is often implied. Directors who agreed on what GenAI could do still diverged substantially on what it meant for their firms and sectors, and this divergence persisted even as they acknowledged capability advances. By capturing these interpretations as they formed and stabilized, we provide an empirical record of how directors reasoned about GenAI during a period of high uncertainty—helping to shed light to the link between technological advancement and competitive expectation.

van Dijk) over 2023–2024 and 2023–2025 (Appendix G). Both exercises are correlational and cannot adjudicate causal effects or the accuracy of expectations; they are offered only as triangulation.

References

- Acemoglu D. 2024. *The Simple Macroeconomics of AI*. National Bureau of Economic Research, Cambridge, MA: w32487. Available at: <http://www.nber.org/papers/w32487.pdf>.
- Agrawal AK, Gans J, Goldfarb A. 2021, May 1. Ai Adoption and System-Wide Change. SSRN Scholarly Paper, Rochester, NY. Available at: <https://papers.ssrn.com/abstract=3847556>.
- Appel R *et al.* 2025, November 19. Anthropic Economic Index report: Uneven geographic and enterprise AI adoption. arXiv. Available at: <http://arxiv.org/abs/2511.15080>.
- Baldwin CY. 2024. *Design Rules, Volume 2: How Technology Shapes Organizations*. MIT Press, 2.
- Baldwin CY, Clark KB. 2000. *Design rules. 1: The power of modularity*. MIT Press: Cambridge, Mass.
- BCG *et al.* 2024, January 8. BCG AI Radar: From Potential to Profit with GenAI. *BCG Global*. Available at: <https://www.bcg.com/publications/2024/from-potential-to-profit-with-genai> [20 January 2026].
- Benner MJ, Tushman M. 2002. Process Management and Technological Innovation: A Longitudinal Study of the Photography and Paint Industries. *Administrative Science Quarterly* **47**(4): 676–707.
- Benner MJ, Waldfoegel J. 2016. The Song Remains the Same? Technological Change and Positioning in the Recorded Music Industry. *Strategy Science* **1**(3): 129–147.
- Bick A, Blandin A, Deming DJ. 2026. The Rapid Adoption of Generative AI. *Management Science* : mns.2025.02523.
- Bresnahan T. 2024. What innovation paths for AI to become a GPT? *Journal of Economics & Management Strategy* **33**(2): 305–316.
- Bresnahan TF, Trajtenberg M. 1995. General purpose technologies 'Engines of growth'? *Journal of Econometrics* **65**(1): 83–108.
- Brynjolfsson E, Chandar B, Chen R. 2025a. Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.
- Brynjolfsson E, Li D, Raymond L. 2025b. Generative AI at Work*. *The Quarterly Journal of Economics* : qjae044.
- Brynjolfsson E, Mitchell T, Rock D. 2018. What Can Machines Learn and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings* **108**: 43–47.
- Chatterji A *et al.* 2025. How People Use ChatGPT. National Bureau of Economic Research.
- Choudary SP. 2025. *Reshuffle: Who wins when AI restacks the knowledge economy*, First edition. Amazon Fulfillment Poland: Wroclaw.
- Choudhury P, Moeen M, Wormald A. 2025. Abductive Empirical Studies: Cultivating and Embracing Plurality in Strategic Management Research.
- Csaszar F, Jia N. 2026. *Handbook of Artificial Intelligence and Strategy*. EDWARD ELGAR PUBLISHING: S.I.
- Daniotti S, Wachs J, Feng X, Neffke F. 2026. Who is using AI to code? Global diffusion and impact of generative AI. *Science* : eadz9311.
- Dell'Acqua F *et al.* 2023. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *SSRN Electronic Journal*. Available at: <https://www.ssrn.com/abstract=4573321>.

- Doshi AR, Bell JJ, Mirzayev E, Vanneste BS. 2024. Generative artificial intelligence and evaluating strategic decisions. *Strategic Management Journal* n/a(n/a). Available at: <https://onlinelibrary.wiley.com/doi/abs/10.1002/smj.3677>.
- Dutton JE, Jackson SE. 1987. Categorizing Strategic Issues: Links to Organizational Action. *The Academy of Management Review* 12(1): 76.
- Eggers JP, Kaplan S. 2009. Cognition and Renewal: Comparing CEO and Organizational Effects on Incumbent Adaptation to Technical Change. *Organization Science* 20(2): 461–477.
- Eggers JP, Park KF. 2018. Incumbent Adaptation to Technological Change: The Past, Present, and Future of Research on Heterogeneous Incumbent Response. *Academy of Management Annals* 12(1): 357–389.
- Eloundou T, Manning S, Mishkin P, Rock D. 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science* 384(6702): 1306–1308.
- Ethiraj SK, Levinthal D, Roy RR. 2008. The Dual Role of Modularity: Innovation and Imitation. *Management Science* 54(5): 939–955.
- Felten EW, Raj M, Seamans R. 2023, April 10. Occupational Heterogeneity in Exposure to Generative AI. SSRN Scholarly Paper, Rochester, NY.
- Flick U. 1998. *An introduction to qualitative research*, 5. ed. Sage: Los Angeles, Calif.
- Gans J. 2025. *The microeconomics of artificial intelligence*. The MIT Press: Cambridge, Massachusetts.
- Gilbert CG. 2005. Unbundling the Structure of Inertia: Resource Versus Routine Rigidity. *Academy of Management Journal*. Academy of Management 48(5): 741–763.
- Gilbert CG. 2006. Change in the presence of residual fit: Can competing frames coexist? *Organization Science*. Institute for Operations Research & the Management Sciences (INFORMS): US 17(1): 150–167.
- Handa K et al. 2025, February 11. Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations. arXiv.
- Hartley J, Jolevski F, Melo V, Moore B. 2024, December 18. The Labor Market Effects of Generative Artificial Intelligence. SSRN Scholarly Paper, Social Science Research Network: Rochester, NY. Available at: <https://papers.ssrn.com/abstract=5136877>.
- Henderson RM, Clark KB. 1990. Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms. *Administrative Science Quarterly* 35(1): 9.
- Hoffmann M, Boysel S, Nagle F, Peng S, Xu K. 2025, April 18. Generative AI and the Nature of Work. SSRN Scholarly Paper, Social Science Research Network: Rochester, NY.
- Iansiti M, Lakhani KR. 2020. Competing in the age of AI: strategy and leadership when algorithms and networks run the world. Harvard Business Review Press: Boston, Massachusetts.
- Jackson SE, Dutton JE. 1988. Discerning threats and opportunities. *Administrative Science Quarterly*. The Johnson Graduate School, Cornell University: US 33(3): 370–387.
- Jacobides MG, Brusoni S, Candelon F. 2021. The Evolutionary Dynamics of the Artificial Intelligence Ecosystem. *Strategy Science* 6(4): 412–435.
- Johnston AC, Makridis C. 2025, July 31. The Labor Market Effects of Generative AI: A Difference-in-Differences Analysis of AI Exposure. SSRN, Rochester, NY.
- Kaplan S, Tripsas M. 2008. Thinking about technology: Applying a cognitive lens to technical change. *Research Policy*. Elsevier 37(5): 790–805.

- Makridis C. 2025, August 21. Organizational Transmission of AI: Managerial Influence on Generative AI Adoption. SSRN Scholarly Paper, Social Science Research Network: Rochester, NY. Available at: <https://papers.ssrn.com/abstract=5401053>.
- Manning SJ, Aguirre T. 2026. How Adaptable are American Workers to AI-Induced Job Displacement? National Bureau of Economic Research.
- Martinez J. 2025. Putty-Clay Automation. Available at: <https://josebamartinez.com/pdf/PuttyClayAutomation.pdf>.
- McAfee A. 2024. The Economic Impact of Generative AI. MIT Sloane.
- Mitchell W, Wu Z, Bruton GD, Gautam DK. 2022. Microlevel Analysis of Institutional Intermediation in a Rudimentary Market-Based Economy: Entrepreneurship in Kathmandu's Indrachok Market. *Organization Science* **33**(6): 2106–2134.
- Noy S, Zhang W. 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science* **381**(6654): 187–192.
- Peng S, Kalliamvakou E, Cihon P, Demirer M. 2023, February 13. The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. arXiv.
- Puranam P. 2025. Re-Humanize: How to Build Human-Centric Organizations in the Age of Algorithms. Penguin Random House SEA.
- Rogers EM. 2003. *Diffusion of innovations*, 5th ed. Free Press: New York.
- Shane S. 2000. Prior Knowledge and the Discovery of Entrepreneurial Opportunities. *Organization Science* **11**(4): 448–469.
- Staw BM, Sandelands LE, Dutton JE. 1981. Threat-rigidity effects in organizational behavior: A multilevel analysis. *Administrative Science Quarterly*. The Johnson Graduate School, Cornell University: US **26**(4): 501–524.
- Suleyman M, Bhaskar M. 2023. The coming wave: AI, power and the twenty-first century's greatest dilemma. The Bodley Head: London.
- Svanberg M, Li W, Fleming M, Goehring B, Thompson N. 2024. Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision? *SSRN Electronic Journal*.
- Teece DJ. 1986. Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy* **15**(6): 285–305.
- Timmermans S, Tavory I. 2012. *Abductive Analysis: Theorizing Qualitative Research*. University of Chicago Press: Chicago, IL.
- Tomlinson K, Jaffe S, Wang W, Counts S, Suri S. 2025, December 22. Working with AI: Measuring the Applicability of Generative AI to Occupations. arXiv. Available at: <http://arxiv.org/abs/2507.07935>.
- Tripsas M. 1997. Unraveling the Process of Creative Destruction: Complementary Assets and Incumbent Survival in the Typesetter Industry. *Strategic Management Journal*. Wiley **18**.
- Tripsas M, Gavetti G. 2000. Capabilities, cognition, and inertia: evidence from digital imaging. *Strategic Management Journal* **21**(10–11): 1147–1161.
- Wu B, Wan Z, Levinthal DA. 2014. Complementary assets as pipes and prisms: Innovation incentives and trajectory choices. *Strategic Management Journal* **35**(9): 1257–1278.
- Yiu S, Seamans R, Raj M, Liu T. 2025, March 20. Strategic Responses to Technological Change: Evidence from Online Labor Markets. SSRN Scholarly Paper, Rochester, NY.

Appendix

Pg Contents

43	<u>Appendix A: Stage 1 – Open-Ended Roundtables (Feb 2024)</u>
43	<u>A1. Respondent Recruitment & Roundtable Composition</u>
45	<u>A2. Insights from Roundtables</u>
53	<u>Appendix B: Stage 2 – Survey (Feb-Mar 2024)</u>
53	<u>B1. Survey & Data Collection</u>
53	<u>B2. Descriptive Patterns</u>
55	<u>B3. OLS Regressions</u>
59	<u>Appendix C: Stage 3 – Structured Roundtables (Apr-Jul 2024)</u>
59	<u>C1. Respondent Recruitment & Roundtable Composition</u>
64	<u>C2. Qualitative Coding & Exemplars</u>
71	<u>C3. Further Details into Investigation</u>
81	<u>Appendix D: Stage 4 – Survey, One-to-One Interviews, Roundtable (Apr-Jul 2025)</u>
81	<u>D1. Survey & Data Collection</u>
85	<u>D2. One-to-One Interviews</u>
88	<u>D3. External Roundtable</u>
92	<u>Appendix E: Stage 5 – One-to-One Interviews (Jan-Mar 2026)</u>
92	<u>E1. One-to-One Interviews</u>
100	<u>E2. External Roundtables</u>
103	<u>Appendix F: Additional Analysis –Matching Stage 2 Survey to Wayback Machine</u>
103	<u>F1. Methodology & Data Collection</u>
106	<u>F2. Descriptive Patterns</u>
107	<u>F3. OLS Regressions</u>
113	<u>Appendix G: Additional Analysis – Matching Stage 2 Survey to FAME Firm Data</u>
113	<u>G1. Methodology & Data Collection</u>
116	<u>G2. Descriptive Patterns</u>
117	<u>G3. OLS Regressions</u>

Appendix A

Stage 1 Open-Ended Roundtables (Feb 2024)

A1. Respondent Recruitment & Roundtable Composition

The design for the abductive stages of our investigation began in December 2023, when our collaboration with the Institute of Directors (IoD) was formally approved and funded through the UK's Regional Innovation Fund. The terms of the collaboration afforded broad methodological and design freedom: our inquiry could take any form and employ any methods, provided the study addressed the implications of GenAI for the UK economy. This latitude was essential for the abductive logic underpinning our research design, as it enabled us to structure each stage in response to emergent findings rather than committing to a fixed protocol at the outset.

In January 2024, we designed Stage 1 in consultation with and subsequent approval from the IoD. We elected to conduct open-ended roundtables, broadly grouped by industry, as the initial mode of inquiry. This choice reflected two considerations. First, given our interest in understanding whether and how heterogeneity in competitive expectations developed, and the absence of a priori hypotheses to test, industry-specific roundtables offered the most informative design: they allowed us to observe where executives within the same sub-sector naturally agreed and disagreed, capturing deliberative dynamics that would be diluted in mixed-industry settings where discussions would likely remain at a more general level (Flick, 1998; Timmermans and Tavory, 2012). Second, the choice of industry groupings was informed by the state of GenAI diffusion at the time. In early 2024, much of the discussion surrounding GenAI (and broader discussion on AI) centered on knowledge work specifically, implicating many of the core tasks and activities performed by knowledge-intensive professional services firms (Brynjolfsson and Mitchell, 2017; Eloundou *et al.*, 2024; Felten, Raj, and Seamans, 2021).

We therefore organized three industry-specific roundtables around sectors with high exposure to knowledge-work applications: non-regulated professional services (management consulting, technology consulting), regulated professional services (legal services, financial services), and media and communications. The distinction between regulated and non-regulated professional services allowed us to separate legal and financial advisory firms, where professional liability and certification requirements structure competitive dynamics, from management consulting and technology consulting, where such institutional constraints are less binding. Media and communications constituted a third natural grouping, as the multi-modal capabilities of GenAI models prevalent at the time (e.g., Midjourney v6, Dall-E 3) were particularly salient for content-intensive work in advertising, and media production.

Alongside the three industry-specific roundtables, we scheduled two omnibus sessions. These served a dual purpose: to examine whether patterns observed within specific industries also manifested at a higher level of generality, and to accommodate sub-sectors that would not be represented in the three industry-focused sessions. As the table below indicates, the omnibus

roundtables captured a wider array of sectors, including education, venture capital, construction, and EdTech, among others.

Recruitment proceeded through established IoD channels. In January 2024, our IoD liaison distributed an email to members of the IoD's policy advisory groups, explaining that this was a collaboration between a leading business school (named in the email) and the IoD to explore the implications of GenAI for firms and sectors. The email presented the five roundtable sessions, described each session's thematic focus, and invited recipients to self-select into the roundtable most relevant to their professional context. No additional screening criteria were applied beyond the requirement that participants hold director-level or equivalent positions; a condition ensured by the IoD's membership structure.

The roundtables were conducted in person at the collaborating institution's campus in February 2024, with each session lasting approximately two and a half hours. The research team facilitated all sessions, posing open-ended questions about participants' experiences with and expectations regarding GenAI, and encouraging participants to share perspectives, respond to and challenge one another, and keep their contributions concrete and context specific. All sessions were audio-recorded with participants' informed consent and subsequently transcribed for analysis. Participants were assured that their responses would be anonymized, and that no individually identifiable information would appear in any published output.

The table below summarizes the participant composition across the five roundtables. Of the 102 executives who registered, 96 attended, yielding a non-attendance rate of approximately 6%. The resulting sample comprised senior executives spanning a range of knowledge-intensive sectors, providing the breadth of perspectives necessary for the exploratory aims of this stage.

#	Theme	Participants	Sector Coverage
1	Non-regulated professional services	19	Management consulting, technology consulting, engineering and construction consulting, healthcare and life sciences consulting, education consulting
2	Omnibus (1)	20	Education, management consulting, finance and risk management, technology and energy, data analytics, professional training, software development
3	Media & communications	19	Communications, PR and marketing, media production, advertising, market research, business services
4	Regulated professional services	16	Legal services (corporate law, litigation, patent and trademark law, immigration), insurance, banking, management consulting
5	Omnibus (2)	22	Management consulting, technology consulting, financial services, banking, engineering and construction consulting, pharma, life sciences

A2. Insights from Roundtables

All five members of the research team were present at each of the five roundtables, ensuring consistent facilitation and enabling real-time triangulation of observations across sessions. All sessions were video recorded and researcher notes were kept digitally. Each roundtable opened with an overview of the research project—its aims, the collaboration with the IoD, and the confidentiality and anonymization protocols governing participation. Following this introduction, we adopted a deliberately open-ended facilitation approach: we began by inviting participants to discuss the broader landscape of GenAI as it related to their firms and sectors, posing initial questions about where the technology had been adopted, where its effects had been felt, and what experiences participants had encountered to date. We then allowed discussion to flow freely, following the directions that participants themselves introduced, in order to observe which topics and concerns were most salient for directors without imposing the research team's prior categories.

At the same time, the research team had discussed and planned a structured probing protocol in advance. When discussion converged on a particular subject—whether a specific application, a competitive concern, or a sectoral dynamic—the facilitators ensured that other members within the sub-sector grouping had the opportunity to respond. The aim was to check for alignment and surface disagreement: to determine whether the perspective being articulated reflected a shared assessment or a more idiosyncratic view, and to encourage participants to challenge, qualify, or elaborate on one another's positions. This combination of open-ended exploration with targeted probing allowed discussions to develop organically while still ensuring that we captured the range of perspectives within each roundtable. The subsections below summarize the substantive themes that emerged across the five sessions.

GenAI, Tasks, & Productivity

Given that our roundtables began with a broad question about GenAI's impact, it is notable that across all five sessions, the discussion immediately gravitated toward the technology's task-level capabilities. Directors spoke first and most readily about their individual workflows: email composition, memo drafting, report summaries, contract review, and pitch deck preparation. The framing was overwhelmingly personal and operational. As one management consultant observed, "right now, it's all task-oriented," noting that the businesses he advised were focused almost exclusively on "the productivity side" and "savings, how they're going to be achieved." A risk management consultant described testing GenAI in her specialist domain and finding the output "pretty decent," adding: "I have an enterprise risk management policy that is decent, but it's very blunt. No way in hell I'm going to present that to a client, but I can spend two hours or six hours or whatever tailoring it rather than punching from scratch." Several directors reported experimenting with the technology for more sophisticated applications and finding it sufficiently capable to serve as a useful starting point, even if outputs required substantial human refinement. One director recounted how a senior developer used an AI coding assistant to identify and resolve software bugs, a process that "would have taken us weeks to find" but was completed "in less than a day," though "it still required the skill set of a high-level developer to find and check that code." A

research director at a management consultancy described systematic experimentation with GenAI as a "co-thinker" for brainstorming and strategic analysis, reporting that the team "were also surprised by how deep you can go if you engage in this dialogue, in this conversation with the machine."

Upon probing the implications of these capabilities, directors in professional services, financial services, and media and communications were particularly vocal about what they perceived as a consequential downstream concern: the development of junior talent. The tasks that GenAI performed most capably—research synthesis, first-draft production, data structuring, routine analysis—were precisely the tasks through which junior professionals traditionally built expertise. Directors recognized this tension readily and articulated it with considerable force. A consultant with experience advising large professional services firms put the concern starkly: "How are we going to use all these juniors consultants who used to take eight days to do a task they can do now in not five minutes, but half a day." An education services director framed the issue in terms of the underlying learning process: "It is the process of researching and structuring and coming up with the answers, not the knowledge at the end. And if AI is giving the student the knowledge instantly, how do they then develop those skills which the workplace needs, such as critical thinking, analysis and so forth?" This concern surfaced independently across multiple roundtables and consistently elicited visible agreement—nods, follow-up examples, and elaboration from other participants—suggesting a widely shared recognition that GenAI's productivity gains at the task level carried implications for how tacit knowledge is accumulated within organizations.

The discussion of task-level capabilities also raised questions about signaling and the assessment of candidate quality. Directors in professional services—financial advisory, consulting, and legal services in particular—noted that traditional hiring instruments such as cover letters, CVs, and written assessments functioned as proxies for a combination of analytical capability, attention to detail, and motivation. A careers consultant specializing in finance and fintech observed that candidates were already using GenAI to structure their applications, but that the results often "just makes it worse" because candidates lacked the skill to use the technology effectively. The broader concern, however, was what happened when candidates could use it well: if application materials uniformly appeared polished and closely matched job descriptions, the informational content of these signals would erode. As one participant noted, where previously "you would take a nice, polished outcome" and "from that would infer some combination of native intelligence, ability to understand the task and motivation to do the work," GenAI severed this inference—"something that looks polished" could now be produced "with the press of a button." This concern extended beyond external recruitment to internal evaluation and advancement. A head of engineering at a technology consultancy posed the question directly: "How do we assess me working in software development, working with software engineers, who now can use ChatGPT, GitHub Copilot to generate code. How do we go about assessing that?" An executive coach working with multinational firms captured the breadth of the challenge: "I think it needs a paradigm shift in what we educate for, what we hire for, what we promote for."

Taken together, these discussions revealed a consistent pattern in how directors engaged with GenAI's implications during the initial roundtables. Directors were readily able to connect the technology's capabilities—capabilities they had personally tested and found substantively useful—to concrete processes within their firms. Yet the connections they drew were overwhelmingly oriented toward productivity and occupational implications: efficiency gains, workforce composition, talent development, and assessment procedures. The competitive and strategic dimensions of GenAI—how it might reshape competitive advantage, alter industry structure, or shift the distribution of value across firms—emerged as a distinct and later conversation, one that did not follow straightforwardly from the task-level assessments that dominated early discussion.

GenAI Adoption Across Sectors

Our roundtables revealed varying degrees of GenAI individual and firm-level adoption across sectors. Initially, we encouraged participants to speak about their own sectors, rather than abstracting to others, and to focus what is *currently happening*, rather than the future.

Our analysis identified two primary adoption patterns: top-down and bottom-up. Regulated professional services, particularly law firms, reported implementing GenAI solutions through top-down initiatives, from developing in-house models and conducting pilot studies to deploying AI solutions at localized scales and engaging with third-party, off-the-shelf solutions. Legal executives spoke of the technology's ease of deployment and its ability to complement existing practices—for instance, by reducing time spent on research or enhancing efficiency in contract review. Notably, legal executives described the technology primarily in terms of complementary and iterative interaction between employees and AI, emphasizing how it augmented rather than replaced human expertise.

In contrast, non-regulated professional services and media and communications sectors described a more bottom-up approach. While few firms in these sectors had established formal top-down policies on GenAI, they did acknowledge informal use. Executives from management consultancies mentioned using AI to provide initial directions for client advisory, assisting in knowledge synthesis, and generating first drafts of pitch decks. Again, these outputs served as a starting point for further refinement and analysis by humans.

Executives in media and communications, particularly PR and advertising, also reported bottom-up adoption patterns, but tended to discuss AI applications more in terms of displacement rather than complementarity. One media executive observed, “A task which previously would have taken the junior a day to do, now they just put it in ChatGPT and get it done in a few minutes.” Executives expressed concern about their inability to accurately gauge the time or effort required for tasks traditionally performed by junior staff. They admitted that their interaction with AI-generated content was often limited to “proof-checking” and described challenges in distinguishing between the respective outputs of AI and humans. Executives also expressed uncertainty about how to evaluate junior staff performance and skill development in light of AI adoption.

Financial services executives reported the lowest levels of adoption, aside from personal use for tasks such as composing emails. While these executives suspected informal use among employees, they generally advocated for a more conservative approach, often citing security concerns. One banking executive cited prior security leaks that had made the bank exceptionally cautious about adopting new technologies in haste. In contrast, participants from the EdTech sector, many from smaller organizations, reported using the technology more frequently.

Emergence of Three Dimensions of Competitive Implications

As roundtable discussions progressed beyond the initial focus on tasks and productivity, a notable pattern emerged in how directors engaged with GenAI's competitive implications. In three of the five roundtables, the conversation shifted from productivity discussions to questions of firm differentiation—how GenAI might be leveraged to strengthen or reconfigure a firm's competitive position. This transition served as a natural bridge, connecting directors' assessments of the technology's capabilities with their reasoning about what those capabilities meant for firm advantage. In the remaining two roundtables, the discussion moved first to questions of displacement—whether GenAI could substitute for firms' existing offerings or enable competitors and entrants to replicate what incumbents provide. Rather than treating the technology's competitive meaning as a single, undifferentiated question, directors' deliberations in all five sessions organized into separable topics, each eliciting distinct forms of reasoning and distinct patterns of agreement and disagreement. When the research team reviewed the discussion transcripts retrospectively, it became clear that directors were invoking different concepts and advancing separate arguments across what we came to identify as three dimensions of competitive implications: competitive differentiation, firm displacement, and industry disruption. Directors engaged intensively on each dimension—debating, offering examples, and challenging one another's assessments—before the conversation reached apparent saturation and a qualitatively different concern surfaced renewed deliberation.

The first dimension to emerge was **competitive differentiation**—the question of whether and how firms could leverage GenAI to improve or reconfigure their value propositions relative to rivals. That this dimension surfaced earliest was itself revealing: directors' instinct, upon moving beyond task-level discussion, was to reason about advantage rather than threat. Directors frequently articulated this dimension in opportunity terms, focusing on how GenAI could enable firms to occupy more defensible competitive positions. A director of research at a management consultancy argued that while operational applications of GenAI would "become standard" and be "embedded in the Microsoft Office, in Word, in PowerPoint," the strategic use of GenAI as a thinking partner represented a genuine source of advantage: "if you manage to be a master on the co-thinking, let's create a uniqueness, something that you can really have an advantage on." The implication was that GenAI's basic productivity applications would be universally accessible and therefore competitively neutral, whereas more sophisticated uses—strategic reasoning, creative synthesis, complex problem framing—offered a window for differentiation precisely because they demanded capabilities that not all firms possessed. A strategy consultant in the mining and metals sector offered a concrete illustration: clients were already arriving with "extremely pointed" questions

because "90% of the work is done internally," leaving consultants to compete only on "deep triangulation" and specialized insight—a shift he saw as an opportunity for firms positioned at the high end. A consultancy owner specializing in government modernization similarly argued that consulting firms' strategic imperative was to operate at the high-value end of the capability spectrum: "that's where they're valued. That's what matters to them, that they can provide high level of skills, perhaps unique insight to their customers. That's how they win their next business." Directors also identified specific assets that they believed would anchor differentiation in a GenAI-enabled landscape. A research director at a large management consultancy argued that firms' accumulated advantages—"trust, reputation, proprietary knowledge, change management" capabilities—would become more rather than less valuable as GenAI commoditized analytical outputs, because clients would still require the judgment, accountability, and relational depth that these assets embodied. A director in financial services and fintech noted that firms actively using GenAI for product innovation faced a market-based reality check: "if it's not adding value, then very quickly the market will not be interested," suggesting that differentiation through GenAI would be disciplined rather than speculative. An energy sector consultant captured the prevailing sentiment more directly: "it's about harnessing the opportunity. Of course, you've got to manage the threats, whatever that might be. I think if you view it as a threat, you're going to miss the opportunity continuously in terms of what it can bring." These exchanges reflected a widespread but not universal inclination to frame GenAI's competitive meaning in terms of how it could be leveraged, with directors identifying proprietary data, relational capital, and specialized expertise as the foundations on which GenAI-enabled differentiation would rest. Several directors also noted that the very homogeneity of GenAI-produced output might itself create new avenues for differentiation through distinctively human judgment and communication.

The second dimension—**firm displacement**—emerged as a counterpoint to the differentiation narrative, often surfacing when directors confronted the possibility that GenAI might enable imitation, disintermediation, or capability substitution that threatened their firms directly. Where the differentiation discussion had centered on what firms could do with GenAI, the displacement discussion centered on what GenAI could do to firms. A management consultant noted that "companies like ourselves, consultancies, for example, where the models are standard," faced a direct threat: "if you have standard models, AI is going to take over, because it's basically a task." The logic here was that firms whose value propositions rested on codifiable, repeatable processes were precisely those most vulnerable to substitution by GenAI-enabled competitors or by clients themselves. The CEO of a neuroscience biotech company drew an explicit historical parallel, likening the current moment to the shift from 2D to 3D computer-aided design: "there's going to be those companies that utilize AI and there's going to be those companies that go bust. Which do you want to be?" A co-founder of an engineering firm framed the competitive pressure more starkly still: "if you are willing, you're more accepting of risk as an organization. So if you're an organization that sits in the pie of a market, you just have to eat your little slice, you're finished. Because someone else who isn't will kill you at a quicker pace than anything normally happens in the commercial market." Directors also raised the pricing mechanism through which displacement would operate: if GenAI enabled a competitor to deliver in a day what previously took two weeks,

fee structures would come under pressure even where the underlying service remained in demand. One director reported that a major professional services firm was "genuinely looking to take less people in on the middle ground where they are going to use AI," providing concrete evidence that displacement planning was already under way at scale.

Yet within this dimension, directors disagreed substantively, and these disagreements proved to be among the most generative exchanges in the roundtables. Several directors pushed back against the displacement thesis, arguing that the barriers to substitution were higher than the threat narrative implied. A research director at a large management consultancy was skeptical that clients would replace professional services with GenAI tools directly: "I will doubt that tomorrow a buyer will download a bot from the open AI store to do some of the critical stuff that they now buy from consultants." He situated this skepticism in historical context, noting that predictions of consulting disruption had circulated before without materializing and that "I will be a little bit cautious, and I will look at the fundamental drivers to see if, and to which extent this will be disruptive for consulting or just incremental." Other directors offered a more nuanced position, suggesting that displacement would be stratified rather than wholesale—that leading firms with established reputations and deep client relationships would weather the transition while second-tier firms operating on more commoditized service models would face existential pressure. This disagreement within the displacement dimension—between directors who saw existential threat, those who saw barriers to substitution, and those who anticipated selective displacement along the quality spectrum—animated sustained debate across multiple roundtables.

The third dimension—**industry and sectoral disruption**—tended to emerge later in the discussions, after firm-level deliberations had been substantially explored. Where the displacement dimension concerned whether specific firms would survive, the sectoral disruption dimension concerned whether GenAI would reorganize value creation and value capture at the level of entire industries—altering who participates, how value chains are configured, and where profits accrue. The shift from firm-level to sector-level reasoning was often marked by directors broadening their frame of reference from their own firms to the architecture of their industries as a whole. One IoD director framed this in sweeping terms during an omnibus roundtable: "existing ecosystems are all going to get disturbed hugely, and value chains are going to get, some will get decimated, others will change beyond recognition." Another director anticipated a wave of entrepreneurial entry: "there are probably thousands of young entrepreneurs looking at this stuff and thinking, how am I going to disrupt an industry and disrupt a sector?" The language directors employed was itself revealing. When discussing firm-level threats, they tended to reach for terms such as "threat," "take over," "replace," or "finished"; when they spoke about wholesale structural transformation, "disruption" was the term they invoked. An engineering entrepreneur made this scale explicit: "the AI is extremely disruptive, probably more so disruptive than anything we've ever had before. It's like a computer in the internet, but on steroids, it's times 10. So this notion of how we assess students, how we assess our staff on payroll, everything is up for grabs."

Within this dimension, directors attended to two factors that distinguished sector-level dynamics from firm-level ones. The first was regulation—not regulation of AI specifically, but the regulatory

architecture governing particular sectors. The engineering entrepreneur pointed to the writers' and actors' strikes in Hollywood as evidence that where "something could have been highly disruptive and was already on its way to be highly disruptive," collective action at the industry level had attenuated its force. Directors in regulated professional services similarly invoked professional liability, certification requirements, and statutory frameworks as sector-level constraints that would shape the trajectory of disruption regardless of individual firms' strategies. The second factor was inter-sectoral heterogeneity: directors recognized that GenAI's disruptive potential would be unevenly distributed across industries, with one director asking whether "the legal profession is going to be impacted more because it's easy for AI to pull in that data" compared to sectors where data availability was more constrained. This prompted extended discussion about which structural features of an industry—data accessibility, regulatory barriers, the codifiability of core tasks—would determine its vulnerability to restructuring.

Yet the sector-level dimension also generated disagreement. Some directors anticipated wholesale restructuring driven by new entrants and the unbundling of traditionally integrated value propositions. Others were more skeptical, pointing to historical evidence that despite recurring predictions of sector-level disruption, the dominant pattern over the past two decades had been consolidation and the entrenchment of incumbent firms. This tension—between expectations of structural transformation and evidence of structural persistence—recurred across roundtables and informed the research team's decision to examine sector-level dynamics more systematically in subsequent stages.

Across the five roundtables, the transition between dimensions followed a broadly consistent pattern. As noted, three of the five sessions moved from productivity to differentiation first, while two moved to displacement—but in all cases, the three dimensions eventually surfaced as separable objects of deliberation. Directors initially gravitated toward differentiation—identifying how GenAI could be leveraged for advantage. As this discussion matured, displacement concerns surfaced as a counternarrative, often triggered by concrete examples of tasks that GenAI could perform at equivalent quality but lower cost. Once the firm-level debate reached a degree of saturation—with directors having articulated their positions and engaged with counterarguments—the conversation shifted to sectoral-level questions: whether value chains would be restructured, whether new entrants would reshape industry boundaries, and whether regulatory and institutional constraints would buffer or channel these dynamics. Importantly, the arguments directors advanced within each dimension were qualitatively distinct (Charmaz, 2006; Timmermans and Tavory, 2012). Differentiation arguments centered on leveraging capabilities and securing advantage; displacement arguments centered on vulnerability, substitution, and competitive pressure; and disruption arguments centered on structural transformation, regulation, and the reconfiguration of industry architecture. Directors did not merely apply the same reasoning at different scales—they drew on different logics, different evidence, and different concerns when engaging with each dimension. This sequential emergence was not perfectly uniform across every roundtable, but the three-part structure recurred with sufficient consistency to suggest that directors themselves perceived competitive differentiation, firm displacement, and industry

disruption as distinct, if related, aspects of GenAI's competitive implications—a perception that informed the design of the Stage 2 survey instrument.

Appendix B

Stage 2 – Survey (Feb-Mar 2024)

B1. Survey & Data Collection

We administered a Qualtrics survey to director-level executives, focusing on GenAI's competitive implications and organizational strategies. Respondents were drawn from the Institute of Directors (IoD) membership pool, with self-selection and no exclusion criteria beyond the requirement that respondents hold director-level or equivalent positions. To enhance data quality, attention checks were embedded throughout the instrument, and item order was randomized to minimize order effects.

The IoD typically conducts monthly surveys of its membership base, coordinated through the same IoD liaison who facilitated our Stage 1 roundtables. For our study, the IoD distributed the survey as a separate instrument, explicitly distinguishing it from its regular polling. This approach has precedent, as the IoD has previously collaborated with academic researchers on dedicated survey instruments.

Of 486 total responses, 222 were fully completed, representing our main analysis sample and a response rate of approximately 11% from the IoD membership base of approximately 2,000 UK members. We found no systematic differences between partial (greater than 50% complete) and full respondents across key demographic and attitudinal variables, suggesting that attrition was not systematically related to the constructs of interest.

Additionally, we administered the survey to a select group of alumni and Executive Education participants from the collaborating institution, yielding 54 additional fully completed responses. Statistical tests revealed no significant differences between these respondents and the IoD sample on key outcome variables, supporting pooling for descriptive analyses. However, the regression analyses reported below use the IoD sample only to maintain a consistent sampling frame.

The survey targeted executives at director level and above, with 29.60% of respondents being CXO or president of their firm and 45.74% at director level. All respondents confirmed their executive-level status in response to a survey question. The table below reports descriptive statistics from our survey data.

B2. Descriptive Patterns

We summarize the key descriptive patterns from the survey. These patterns served to triangulate qualitative observations from Stage 1 and to motivate the regression analyses in Section B3.

GenAI impact within organizations thus far

When asked to describe GenAI's impact within their organizations to date, the majority of directors reported limited effects: 27.2% reported no impact and 30.4% reported only minor impact, while

20.3% reported moderate impact, 14.3% significant impact, and 7.8% transformative impact. Firm size appeared to shape these assessments. Smaller firms reported more polarized experiences, with both the highest rates of no impact (24.8%) and the highest rates of significant or transformative impact (29.4%), consistent with Stage 1 observations that smaller firms exhibited greater variance in adoption. Organizational modularity emerged as a notable correlate: high-modularity organizations reported significant or transformative impact at more than double the rate of low-modularity organizations (33.9% versus 14.6%), consistent with the view that modular structures facilitate GenAI deployment within discrete organizational components.

Organizational versus industry expectations

Directors exhibited a consistent discrepancy between their expectations for GenAI's impact at the organizational and industry levels. At the organizational level, 84.8% of executives anticipated at least some meaningful change, with 40.4% expecting operational efficiency gains, 23.3% foreseeing competitive transformation, and 21.1% anticipating business-model reinvention. At the industry level, expectations were more pronounced: 89% anticipated at least moderately significant impact on their industry over the next five years. This discrepancy—higher industry-level than organizational-level expectations—is consistent with directors perceiving GenAI as a structural force whose sectoral implications exceed what any individual firm will experience internally.

Distinctiveness potential versus competitive imitation

The survey captured directors' assessments along two key competitive dimensions: the potential for GenAI to unlock new forms of distinctiveness, and GenAI's effect on the ease of competitive imitation. These two distributions exhibited markedly different properties. Distinctiveness potential displayed a clear positive skew, with 91.9% of directors perceiving at least some positive potential. In contrast, assessments of competitive imitation approached a more uniform distribution, with substantial minorities anticipating both greater ease (37.7%) and greater difficulty (35.0%) of imitation, and the remainder perceiving little change. These differing distributional properties suggest that the factors shaping perceptions of GenAI-enabled differentiation may be distinct from those shaping perceptions of GenAI-facilitated imitation—a conjecture that the regression analyses examine directly.

Within narrowly defined sub-sectors, directors exhibited marked disagreement on competitive implications. In management consulting ($n = 36$), standard deviations on disintermediation likelihood ($SD = 1.25$) and distinctiveness potential ($SD = 1.15$) approached the overall sample variance, indicating that sub-sector membership explains little of the observed heterogeneity. Banking and investment directors showed greater consensus on competitive outcomes (disintermediation $SD = 1.08$; distinctiveness challenge $SD = 0.74$), while software and technology directors exhibited the highest variance (distinctiveness potential $SD = 1.32$). Across all outcome variables, 74–100% of total variance occurred within sub-sectors rather than between them, confirming that the heterogeneity observed in Stage 1 roundtables is not reducible to sectoral differences.

Cross-sectoral comparisons corroborated Stage 1's observation that regulated sectors exhibit greater consensus. Directors in highly regulated industries showed markedly lower variance on distinctiveness potential (SD = 1.20 versus 1.87) and GenAI's expected industry importance (SD = 0.94 versus 1.27). Banking and insurance directors converged on competitive expectations; management consulting and software directors exhibited wider dispersion, with standard deviations on distinctiveness potential approximately 1.4 times larger. Notably, we do not find evidence that these differences reflect variation in productivity evaluations—assessments of GenAI's potential for productivity enhancement exhibited similar variance across regulation levels.

B3. OLS Regressions

We estimate ordinary least squares (OLS) regression models with heteroskedasticity-robust standard errors (HC3) to identify systematic associations between organizational characteristics and directors' competitive expectations. While our outcome variables are measured on ordinal scales (1–5 or 1–7), OLS provides easily interpretable coefficients and yields consistent estimates when scale categories are approximately evenly spaced—a standard approach in management research using Likert-type items. We report three specifications for each outcome: (1) bivariate associations, (2) a multivariate model including all predictor and control variables, and (3) the multivariate model augmented with sector fixed effects to absorb unobserved sector-level heterogeneity. Our predictor variables capture organizational characteristics that emerged as salient during Stage 1 roundtables: proprietary data importance, the role of tacit knowledge, technical knowledge, the importance of pattern recognition to the firm's core offerings, the perceived modularity of operations, and the level of industry regulation. We additionally control for relationship knowledge, market insight, certification importance, referrals, and firm size (annual turnover).

Core results

Table 3 in the main text presents OLS results for four core outcome variables that map onto the three dimensions of competitive implications identified in Stage 1: industry importance (industry-level expectations), distinctiveness potential (competitive differentiation), disintermediation likelihood (firm displacement), and financial viability risk (firm displacement). We elaborate on key patterns here.

No single organizational characteristic predicts all dimensions of competitive expectations; rather, different characteristics associate with different dimensions. This finding is consistent with the qualitative observation from Stage 1 that directors invoked distinct reasoning when discussing differentiation, displacement, and industry-level implications.

Pattern recognition importance emerges as the most consistent predictor across specifications. Directors who rate pattern recognition as central to their firm's success perceive lower distinctiveness potential ($b = -0.22$, $p = 0.016$), lower disintermediation likelihood ($b = -0.20$, $p = 0.021$), and lower financial viability risk ($b = -0.17$, $p = 0.040$). In substantive terms, a one-unit increase in pattern recognition importance is associated with approximately one-fifth of a standard

deviation decrease in each of these outcomes. These associations are negative across all three firm-level competitive dimensions, indicating that directors whose offerings rely heavily on pattern recognition perceive GenAI as less consequential for competitive dynamics—a finding we explore further in the extended results below and in Stage 3.

Modularity and regulation each predict a single dimension. Modularity associates negatively with perceived industry importance of GenAI ($b = -0.24$, $p = 0.039$); a one-unit increase in modularity corresponds to a 0.24-point decrease on the five-point industry importance scale, an effect size comparable to one-fifth of a standard deviation. Directors in more modular organizations see GenAI as less transformative for their industries. Regulation associates positively with distinctiveness potential ($b = 0.32$, $p = 0.037$), consistent with the interpretation that directors in regulated industries perceive GenAI as an opportunity to differentiate within protected boundaries—a pattern that aligns with Stage 1 observations from the regulated professional services roundtable, where directors exhibited greater consensus on sectoral stability.

Proprietary data: a confounding note

Proprietary data importance does not reach conventional significance thresholds in the full specification for any of the four core outcomes, despite exhibiting significant bivariate associations with all four. This non-significance does not reflect insufficient variation: proprietary data importance exhibits a standard deviation of 1.21 across the sample, the highest coefficient of variation among the core predictors, with responses spanning the full 1–5 range. Rather, the attenuation from bivariate to multivariate estimates indicates confounding: proprietary data importance correlates with several other predictors in the model—certification ($r = 0.37$), pattern recognition ($r = 0.32$), market insight ($r = -0.32$), regulation ($r = -0.29$)—and bivariate coefficients attenuate 44–75% when these correlated characteristics are included. The apparent bivariate associations between proprietary data and competitive expectations are thus partially attributable to shared variance with other organizational characteristics. We return to this issue in Stage 3, where the structured roundtable format allows us to probe directors' reasoning about proprietary data more directly and to disentangle it from related organizational attributes.

Extended results

To examine whether these patterns extend beyond the four core outcomes, we estimate the same specification (all predictors with sector fixed effects) for four additional outcomes: challenges to maintaining distinctiveness, training challenges, employee productivity enhancement, and the potential to build new products. These outcomes capture the human capital dimension, productivity perceptions, and a broader measure of competitive defensibility. The table below presents these results.

The extended results reinforce and extend the patterns observed in the core specifications. Pattern recognition importance remains the most consistent predictor, associating negatively with three of the four additional outcomes: challenges to maintaining distinctiveness ($b = -0.17$, $p = 0.017$), training challenges ($b = -0.21$, $p = 0.015$), and the potential to build new products ($b = -0.30$, $p =$

0.000). The effect sizes are substantively comparable across these outcomes, with a one-unit increase in pattern recognition importance corresponding to approximately 0.15–0.30 points on each five-point scale. Across all eight outcomes examined, pattern recognition is the only predictor with a consistently negative coefficient, suggesting that directors whose firms rely on pattern recognition perceive lower competitive implications across multiple dimensions.

Two additional findings emerge. First, tacit knowledge importance associates positively with training challenges ($b = -0.32$, $p = 0.005$), its strongest association across all outcomes examined. A one-unit increase in tacit knowledge importance corresponds to a 0.32-point increase in perceived training challenges—an effect comparable in magnitude to approximately one-quarter of a standard deviation. Directors who rate tacit knowledge as central to their firm’s success perceive substantially greater risk that GenAI will challenge talent training, consistent with the concern voiced in Stage 1 roundtables that the tasks GenAI automates are precisely those through which junior professionals develop expertise. Proprietary data importance, by contrast, does not associate with any of the extended outcomes at conventional thresholds, reinforcing the confounding pattern documented above.

Second, the enhancement and productivity outcomes exhibit the strongest overall model fit (R^2 values of 0.34–0.36). Modularity associates negatively with both employee productivity ($b = -0.21$, $p = 0.056$) and the potential to build new products ($b = -0.29$, $p = 0.054$), while pattern recognition associates negatively with the potential to build new products ($b = -0.30$, $p = 0.000$). Directors in modular organizations and those relying on pattern recognition thus perceive less enhancement potential from GenAI—a pattern that parallels the competitive dimensions and suggests a degree of coherence in how directors’ organizational characterizations shape their expectations across both productivity and competitive domains. Regulation associates positively with the potential to build new products ($b = 0.31$, $p = 0.009$), further reinforcing the pattern that directors in regulated industries perceive GenAI as an opportunity within protected boundaries.

Extended Results: Additional Outcome Variables (Sector Fixed Effects)

	Distinctiveness Challenge (e)	Training Challenges (f)	Employee Productivity (g)	Build New Products (h)
	(3)	(3)	(3)	(3)
Proprietary Data	-0.104 (0.258)	0.024 (0.780)	0.005 (0.951)	-0.051 (0.598)
Tacit Knowledge	-0.078 (0.533)	-0.317 (0.005)	-0.199 (0.087)	-0.182 (0.197)
Technical Knowledge	0.040 (0.686)	0.195 (0.082)	-0.051 (0.641)	0.040 (0.730)
Pattern Recognition	-0.166 (0.017)	-0.212 (0.015)	-0.177 (0.024)	-0.302 (0.000)
Modularity	0.001 (0.995)	-0.063 (0.590)	-0.212 (0.056)	-0.291 (0.054)
Regulation	-0.042 (0.670)	0.099 (0.332)	0.140 (0.118)	0.307 (0.009)
Relationship Knowledge	-0.148 (0.205)	0.065 (0.602)	0.159 (0.158)	0.044 (0.770)
Market Insight	0.089 (0.424)	-0.018 (0.871)	0.249 (0.031)	0.448 (0.000)
Certification	-0.206 (0.014)	-0.080 (0.256)	-0.056 (0.446)	0.030 (0.708)
Referrals	0.057 (0.607)	0.050 (0.534)	0.061 (0.484)	-0.064 (0.546)
Firm Size (Turnover)	0.051 (0.386)	0.101 (0.067)	0.067 (0.271)	-0.067 (0.338)
Sector Fixed Effects	Yes	Yes	Yes	Yes
R ²	0.313	0.232	0.337	0.364
Observations	221	221	221	221

Notes: OLS coefficients with exact p-values in parentheses; robust (HC3) standard errors used throughout. All models include all predictor and control variables plus sector fixed effects. Dependent variables: (e) “Rate the likelihood to which you expect GenAI to challenge maintaining distinctiveness in the marketplace within your organisation over the next 5 years” (1–5 scale). (f) “Rate the likelihood to which you expect GenAI to challenge training talent within your organisation over the next 5 years” (1–5 scale). (g) “Which aspects of your organisation stand to benefit from GenAI: Employee productivity” (1–5 scale). (h) “Which aspects of your organisation stand to benefit from GenAI: Build new products/services” (1–5 scale).

Appendix C

Stage 3 – Structured Roundtables (Apr–Jul 2024)

C1. Respondent Recruitment & Roundtable Composition

This appendix provides detailed supporting information for Stage 3 of our study, which consisted of structured roundtables designed to probe the mechanisms underlying directors' competitive expectations regarding GenAI. Building on patterns identified in Stage 2 survey responses, Stage 3 examined how executives' understanding of their firms' value propositions shaped their perceptions of displacement threat, potential to leverage GenAI for competitive differentiation, and expectations of broader industry disruption. The roundtable discussions and participant quotations captured in this stage reveal the interpretive mechanisms through which executives translated task-level capabilities into firm- and industry-level expectations.

We conducted nine roundtables between April and July 2024. Seven were paired roundtables (A1a–A4) constructed using Stage 2 survey responses to group participants according to theoretically salient variables. These variables included perceived importance of pattern recognition, proprietary data, competitive uncertainty, and modularity. Directors were matched into sessions based on high or low scores on these constructs, allowing us to compare how executives with similar structural perceptions reasoned about GenAI's implications. This design inverted the logic of Stage 1, where directors had been grouped by sector; in Stage 3, participants from similar sectors were intentionally distributed across sessions to observe whether similar reasoning emerged independently of direct peer confrontation.

The composition of the nine Stage 3 roundtables is summarized below.

#	Theme	Participants	Sector Coverage
A1a	High pattern-recognition importance	23	Pharmaceuticals, biotechnology, cybersecurity, insurance, management consulting, ESG assessments, behavioral science, environmental planning and landscape design
A1b	Low pattern-recognition importance	13	Management consulting, venture capital, software development, environmental services, alternative dispute resolution, AI ethics certification, digital agency services
A2a	High proprietary data importance	17	Pharmaceuticals, management consulting, legal services (corporate, litigation), insurance, banking, technology advisory, corporate restructuring
A2b	Low proprietary data importance	15	Telecommunications, management consulting, venture capital, chemical manufacturing, marketing, medical services, construction, education
A3a	High competitive uncertainty	11	Management consulting, digital banking, technology solutions, risk and compliance advisory, real estate advisory, manufacturing
A3b	Low competitive uncertainty	18	Management consulting, technology, venture capital, legal services, chemical manufacturing, manufacturing, meteorology, food and beverage export
A4	Modularity (mixed)	15	Healthcare, pharmaceuticals, financial services, technology, management consulting, legal services, construction equipment supply, language translation
B1	Strategy & policy	21	Management consulting, banking, insurance, legal services, pharmaceuticals, telecommunications, ESG assessment, education
B2	Private equity	30	Private equity, venture capital, management consulting, legal services, technology, banking, retail

The first seven roundtables (A1a–A4) focused on variables identified as salient in Stage 2. By pairing high and low configurations of specific constructs, we were able to conduct comparative analysis of directors who differed systematically in how they construed their firms’ assets and competitive environments. This design enabled us to examine divergent interpretations under comparable GenAI scenarios and to identify how specific organizational perceptions shaped expectations of displacement, differentiation, and disruption.

The final two roundtables (B1–B2) were conducted to sense-check the patterns emerging from the A-series sessions and to introduce an outside-in perspective. Whereas the A-series sessions were derived from the IoD survey population and focused primarily on operating directors, B1 and B2 incorporated participants positioned to evaluate firms across sectors rather than within a single

organizational context. B1 comprised senior strategy and policy actors, while B2 comprised predominantly private equity, venture capital, and investment-adjacent executives. Together, these sessions enabled us to assess whether the interpretive patterns identified among operating directors were echoed, challenged, or reframed by external evaluators observing competitive dynamics across multiple firms and industries.

Both external sessions were conducted under Chatham House Rules to facilitate open discussion; participants consented to the disclosure of their identities for transparency regarding cohort composition. The anonymized composition of external roundtable B1 is provided below.

Role / Position	Organization type	Category
Director of Policy and Corporate Governance	Membership/Governance Body	Policy and Governance
Director General	Membership/Governance Body	Policy and Governance
Chair	Membership/Governance Body	Policy and Governance
Managing Director	Management Consulting	Consulting and Advisory
Senior Partner	Management Consulting	Consulting and Advisory
Founder and Principal	Strategy Consulting	Consulting and Advisory
Vice President	Research and Advisory	Consulting and Advisory
Managing Director	Banking and Financial Services	Financial Services
Partner	Financial Services Consulting	Financial Services
Chief Investment Officer	Investment Advisory	Financial Services
Founder and Director	Risk Consulting	Insurance and Risk
Non-Executive Director	Insurance Advisory	Insurance and Risk
Senior Partner	Intellectual Property Law	Legal Services
Managing Director	Intellectual Property Services	Legal Services
Founder and Managing Director	Pharmaceutical Consultancy	Pharmaceuticals
Chief Executive Officer	Telecommunications	Telecommunications
Head of ESG and Sustainability	ESG Assessment and Advisory	ESG and Sustainability
Managing Director	Change Management Consulting	ESG and Sustainability
Professor and Academic Director	Higher Education	Education
Founder and Director	Business Education and Training	Education
Chapter President	IT Governance and Audit	Education

Note: N = 21. This workshop was conducted under Chatham House Rules. Participants consented to the disclosure of their identities for the purpose of providing transparency regarding the composition of the cohort.

While B1 participants were positioned primarily as strategic advisors and policy-involved actors, B2 participants were largely capital allocators and board-level evaluators, allowing us to compare how different forms of external oversight interpreted GenAI's competitive implications.

The anonymized composition of external roundtable B2 is provided below.

Firm	Role / Position	Organization Type
BCG	Global Vice Chair, Global Advantage Practice	Management Consulting
PwC / Strategy&	Partner; Global Technology Leader	Professional Services
General Atlantic	Vice President, Data Science	Growth Equity
Allen and Overy	Partner	Law Firm
Allen and Overy	Partner	Law Firm
Mastercard	VP of Products; Chief of Staff	Fintech (Payments)
Bank of America	Trading Analyst	Banking, Financial Services
Ocado	Chief Data Officer	Online Retail and Tech
IoD	Chair South, Technology	
Laing O'Rourke	Head of Strategy and Transformation	Construction and Engineering
Cenergy Holdings	CFO	Industrial Holdings (Energy)
Mantu	Executive Vice President	Management Consulting
Mantu	Chief of Staff to CEO; AI Transformation	Management Consulting
Eros Media	Chief Technology and Product Officer	Media and Entertainment
Potamitisvekris	Managing Partner	Law Firm
Provenir	Group CFO	Financial Technology (AI)
Fiserv	SVP, Global Digital Solutions	Financial Technology
GoodRx	Chief Product and Technology Officer	Healthcare Technology
Emovid	CEO and Co-Founder	AI Technology Start-up
Claria Partners	Advisor	Strategy Advisory
Claria Partners	Advisor	Strategy Advisory
Le Livre Ouvert	General Manager	Publishing and Retail
London Metropolitan Uni	Director	Higher Education
WonderWorldScience	Director	Education and Science
Business School	Professor of Entrepreneurship and Innovation	Academia
Business School	Associate Professor of Strategy and Entrepreneurship	Academia
Business School	Senior Advisor; Academic Director	Academia
Business School	Researcher	Academia

Note: N = 28. This workshop was conducted under Chatham House Rules. Participants consented to the disclosure of their identities for the purpose of providing transparency regarding the composition of the cohort. Business schools anonymized for peer review.

Following transcription, we employed qualitative coding methodologies (Charmaz, 2006; Gioia, Corley, and Hamilton, 2013) to analyze the roundtable data. We used open coding to identify first-order concepts from informants' own words (Guba and Lincoln, 1982), then employed axial coding to develop higher-level themes, ultimately aggregating these into broader theoretical dimensions (Gioia and Chittipeddi, 1991). This abductive process allowed us to surface the structural perceptions linking directors' understanding of their firms' value propositions to their expectations regarding GenAI's competitive implications.

C2. Qualitative Coding & Exemplars

1 st Order Concepts	2 nd Order Themes	Aggregate Dimensions
Fear that clients will integrate technical knowledge in-house, rendering technical-based value propositions less competitive	Technical Knowledge	Threat of Firm Displacement
Concerns about new entrants with lack of specialized knowledge entering as lower knowledge barriers		
Regulatory requirements deterring new entrants	Regulatory Intermediation	
Dependence on industry-specific certifications and compliance standards that are difficult to meet		
No concern as business is conducted through trust	Sectoral Norms	
GenAI requires experimentation, yet industry remains conservative		
Industry conservatism will insulate from substitution		
Greater difficulty in determining output quality, relational capital becomes more important	Relational Capital	
Supercharging the firm’s business relationships		
Firm value proposition being about identifying and exploiting patterns, fitting right into GenAI’s capabilities	Pattern Recognition (as Vulnerability)	
Fear of competitors generating similar results with different data		
The firm’s value proposition being vulnerable to being chipped away	Modularity (as Exposure)	
New entrants displacing specific parts of the value chain		
Proprietary data being able to supercharge value offerings	Proprietary Data	Leveraging GenAI for Firm Differentiation
Reservoirs of data and communications that can now be organized to enhance the decision-making process and output quality		
Clean and organized data will be a game-changer		
GenAI will be leveraged best by employees with knowledge that is kept in their minds and habits	Tacit Knowledge	
Since tasks are being automated, tacit knowledge will be more difficult to build up, increasing existing firm advantage in expertise		
Not every organization has the right risk appetite	Flexibility to Experiment	
Some organizations have more tolerance from their stakeholders to experiment with the technology early		
Move towards “uncharted land” - Ability to spot and extrapolate patterns will be supercharged by GenAI	Pattern Recognition (as Scalable Advantage)	
We can scale pattern recognition across far more data than was previously possible		
Modular workflows allow us to plug in GenAI incrementally without overhauling entire systems, reducing implementation risk	Modularity (as Integration Opportunity)	
Greater modularity than competitors will allow first mover advantages to consolidate new market positions		
The way patterns are identified and acted upon in the industry is fundamentally changing	Pattern Recognition (Reconfiguration)	Broader Industry Disruption
GenAI is redefining what counts as pattern recognition expertise across sectors		
Some sectors depend on pattern recognition more than others, creating uneven disruption		
GenAI is blurring the boundaries between previously distinct process modules, allowing reorganization of silos	Modularity (Reconfiguration)	
Organizational processes are too complicated as a system to deploy GenAI effectively in the short term		
Entire value proposition of the sector is being picked apart by combination of new entrants and incumbents from other sectors		

Note. Data structure follows the Gioia methodology (Gioia, Corley, & Hamilton, 2013). Pattern Recognition and Modularity appear across all three aggregate dimensions, reflecting the multivalent nature of these constructs in informants' sensemaking—the same phenomenon was simultaneously invoked as a displacement threat, a basis for differentiation, and a driver of industry disruption. First-order concepts may also capture both the informant's own firm perspective and observations about competitors, as executives frequently recognized that the same dynamics applied to both

Exemplar Coding for 2nd Order Themes

2 nd Order Themes	Exemplar Quotes
Technical Knowledge	“I think in terms of our industry, [...] it’s all about the application of technical knowledge. We’re consultant engineers, so it’s about how we look at data and how we interpret it. I think the big fear, [...] especially for us as an SME, is that larger companies may be able to leverage GenAI [...] and render us useless. [...] It’s how we will use GenAI to provide better information and results that’s going to keep us ahead—but if you stay back, I think you’re dangerously getting overtaken. [...] This illustrates the question: did we have an edge because of our ability [to interpret data]?” “Technical knowledge, I think, is more like a multi-step, everyone will know, everyone will learn, everyone will read. So, I don’t think that’s going to make any difference anymore.” “If your business still relies on technical knowledge, then you’re toast, basically. We don’t see any benefit, frankly, in trying to hang on to technical knowledge nowadays. It is no longer a differentiator and won’t distinguish your business in any way, shape, or form.” “The hierarchy of firms is going to shift in our industry for sure. Those firms with technical expertise and sophisticated digital infrastructure are going to realize their differentiation has been gone [...]. We’re not going to be completely substituted, we [humans] are still needed to manage statistical chaos [...] Differentiation is going to more likely center around who best manages this chaos.” “On the technical knowledge piece, I think one of, and I think this transcends industry sectors, there’s been more and more contracting-out, as it were, to specialist organizations. So, I don’t think any of us probably appreciate what the balance is between in-house versus out-of-house, as it were.”
Regulatory Intermediation	“So that’s a big thing. Now, industry to industry, things will differ. In [the] financial services sector, yes, we are based on it. We will test everything because we are regulated. And there are a couple of bottlenecks. I think [there are] a couple of big walls in front of us. So, one, of course, is the AI regulation itself. How would FSFCA, who regulates us, see it? They have not defined it yet. We reach out to them every other day and they have no idea. We then reach out to the PRA, because everything is divided in these countries. I think there’s not one regulation where we can talk and find things. So, it’s not being divided between FCA, PRA, and ICO.” “I used to be a regulator, so I’m still very much in touch with former colleagues there and with [REGION] government. And both of those bodies are very much trying to encourage the adoption of wider fintech and regtech products in order to help the industry. But obviously, I’d say obviously, I think it’s obvious that they cannot give any sort of definitive view to industry to say, ‘Use this product, we think it’s great.’ Clearly, it wouldn’t be appropriate for an authority to endorse a product like that. But industry doesn’t appear to be, or appears to be nervous of taking on any sort of product without that level of endorsement, which they’re not going to get.” “I think some other consulting firms, [CONSULTANCY], for example, they’re just pushing and saying, ‘Well, the financial services...’ They’re trying to sell. [CONSULTANCY] is busy here trying to get more business. Get that reporting. But they are basically saying that advisory by robots. Robo-advisory and stuff like that.” “Technically, even humans as brokers, we are not allowed to advise. We can recommend or we can provide information. Even not recommend. The words are very clear. Advise for this. And now, not a human, you ask an avatar who is picking up 20 tons of information from the dark web and giving advice that, ‘Hey, you should be buying this and that and that.’ And you buy it. And you lose all your money. Who is responsible? This is a huge issue.” “Regulation sets guardrails, rules, ways of doing things, standards that may be harmonized with that regulation. It’s wonderful, we love it. [...] And what have we got? We’ve got a whole ecosystem and economy growing around compliance, around audits, around conformity. [...] So the barriers to entry are quite high and it sets a moat around.” “If the regulatory was not there, I think the financial services industry would have picked up immediately GenAI, because we’ve been doing a lot looking at structured and unstructured data to find insights.” “I think two things naturally go together because if you operate in a highly regulated industry [...] it’s certainly going to mean that you have processes and sign-offs and what I consider to be things, pieces that you could put into a box and explain to a regulator what are the steps that you go through. I guess that’s also [...] steps that AI can look at and it can find ways of doing it more efficiently and quicker.”

Sectoral Norms	“In construction there is quite a famous [CONSULTING FIRM] survey about construction and digitization and being beaten into last place by agriculture and hunting. [...] From our side, it’s slow in adopting; the benefits are potentially massive because they’re themselves and it’s very transactional.” “A part of my background is construction and civil engineering, and there’s a lot of data about it, and it’s not very well used, it’s not very well shared, it’s not very collaborative, and there’s an enormous opportunity, it has been for at least 40 years for applications.” “The pharmaceutical sector, the life sciences industry in general, is an extremely conservative industry. I moved into life sciences from the financial sector, and it was like stepping back a century in terms of the way things got done. And it’s improving, but it is still extremely sluggish to change.” “The fear factor is still there. As people have said, the [SECTOR] industry is conservative, and without wishing to perpetrate any sort of stereotypes, we’re not complete hicks on the sticks in [REGION], but perhaps we’re a little bit more cautious than other people in the big smoke might be. So there is all of that to contend with as well.” “Traditionally, healthcare is conservative, because you have you’re dealing with human life, basically, and there’s so many regulations they have to adhere to. From a government level as well, they’ve been scrutinized, left, right and center. And if things go wrong, it really goes wrong. And it’s built will be kind of all in the media. So that’s why they’re so cautious. [...] It’s a different way of working. And just remember that, I mean, I’m working in the [INSTITUTION]. A lot of people are still IT illiterate. So we still having to train people how to use Office or Excel or Office 365, they will still struggle with Teams. So that is the level really.” “There’s also, I think, resistance to using AI, perhaps because the culture is not really there, that people are willing to innovate, experiment within their organizations. So that’s a perspective I would have, yes.” “One of the major barriers in there, in the university sector is the fact that there are so many incumbents who are people. So you’ve got a huge invested base in the academic community who would not fully appreciate a lot of their work being seen to be downgraded by the implementation of AI. [...] And I do think that the change, the potential change is enormous, but I think there’s quite an embedded resistance. So I think that will take, it’ll take quite a long time. And it’s like any change project in any organization, in any sector. The transition from where you are to where you need to be involves taking people with you. The fact that you’ve got technology that can do this sort of stuff is interesting, but it will not be the driver of the change.” “[SECTOR] and because it’s so creative, people’s, any organizations or individuals’ respective uses differ quite widely, quite easily. [...] A large part of them is guarding that cultural heritage from strange innovation. In that sense they’re very conservative.”
Relational Capital	“I come from a very knowledge-based, you know, professional service in my life, IT service most of my life. And then, you know, marketing and advertising. Where a lot of your work is very much around, you know, knowledge, data, content. And weaving through lots and lots of data that you have accumulated through years of acquisitions and complexity and all of that. So, all of those dynamics exist. It’s just not a big regulated... I just don’t come from a big regulated financial service industry, right? So, I think in many of those industries, operational efficiency, over time as people are learning, will lead to massive initial transformations because costs will come down significantly and people will be able to use their time on more human-type relationship areas of focus.” “Frankly, people know me. It’s a small island, it’s a small industry. [...] If you have any issues, frankly, I’m just around the corner. I can just pop around and talk to you about it. [...] That brings a face, that brings a level of support [...] they know who I am, they know where I am. It brings a level of comfort and reassurance.” “In terms of tacit knowledge in mental health, it’s more about the relationship you have with your clients, which can’t be—that’s what I meant with you can’t codify that. Yes, you can do certain basics. But really kind of having a deep relationship with your patients to help them isn’t—I don’t think it’s going to happen through AI.” “How fast as consumers, as us, as humans, will trust AI systems. In the financial services sector, we’ve tried this for a long time. We said brokers will go away. It has not been better. [...] I would still go out and look for a person to give me good advice. [...] It is because I don’t trust it. Not because the AI is better than the humans, maybe the humans are the same, but because it’s a person I trust.”
Proprietary Data	“You know, collaborations between pharma companies on developing a vaccine involve a lot of intellectual property, licenses, and terms. To illustrate the point, imagine two law firms who have completed French nuclear energy projects. They have to navigate planning laws, indemnities, liability allocation, supply chain issues, government funding, and other external funding. But you can’t just put that into ChatGPT and get an answer to draft a French nuclear energy contract. It’s impossible because the system

	isn't trained on that specific data, which likely exists only in a few cloud-based files. That data is intellectual property—it's incredibly valuable, high-end, and rare. And we have billions of such valuable data points all over the world." "If I look at the second model of strategy consulting, you know, when I started in Booz Allen in 1995, they had a better knowledge management system than most professional services organizations have now, with the exception of now the other McKinsey, Bain, and BCG. When we compete with McKinsey, Bain and BCG, we know that you guys can pull together a deck of 50, 100 pages on an industry topic in a way that no one else can, in a way that actually we can't because we have... our risk management rules prevent us actually from doing that a lot. So that gives you a competitive advantage that we don't have." "We've been collecting data from thousands of client engagements over decades, but we've barely scratched the surface of what insights we might unlock. With generative AI, we're beginning to see our historical data in a new light. It's not just about what we learned from each project individually anymore. We're excited about the patterns and insights that AI might reveal across our entire body of work. We don't even know what questions to ask yet, but we're certain there's immense value hidden in our data that could revolutionize how we serve our clients." "If I take the world of, you know, financial diligence, which we are probably the world number one in doing, we have done more financial diligence exercises than almost anybody else put together. We know how to do that. We've got the historical data. We've got the metrics and we've got a bench strength of thousands of people who can help effectively train a model to do that." "Where they're real weakness is in the warehouses and the robotics and also in the delivery fleets and the routing. [...] They understand the human body well enough to know where the robot should deliver the article and depending on its weight, size and shape [...] so that they can pick it up at their maximum physical effect. [...] That's proprietary knowledge that I suspect nobody else in the world has got at the moment." "Data itself is a huge issue for our industry. [...] We still can't have a 360 view of our customers. Data is across thousands and hundreds of systems, especially in the insurance sector. [...] None of one organization has one policy system or a claim system. We have hundreds of them. It's very complex. So we don't know what is the true value or where is the data." "The opportunity to actually join forces in industry in a way that you can hold on to your own data, but still put your data through a model that enables you to train that model, but you don't have to expose your own data. And that's where I think the industry collaboration, that's when you will start to see [...] some bigger [changes]."
Tacit Knowledge	"In our sector, tacit knowledge is the intricacies behind building a wall and then plastering it. Expert plasterers know exactly what they're doing, and it's quite difficult to replicate that with a machine, though we can do it with brickwork. There comes a point when perhaps we'll be able to do it with plaster as well... Our expectation is that we have the people with the best understanding of these tacit features and hence will be able to use generative AI better than others." "I mean, we've just sort of issued a free chatbot that we've developed, which isn't full generative AI, but it's the only example I've got. It's the best example that I've got, which we've trained up on the AML, CFT, CPF regime in [REGION]. And this is obviously something that we advise on fairly frequently. We've had colleagues, clients and so on coming to us and saying, 'Well, aren't you putting yourselves out of business?' The answer is, well, frankly, no. If you know where to look in the regime, the answers are there already. Our skill is more in what you're calling [the] tacit knowledge side of things. So, the know-how, more of an interpretive skill, the communication side of things." "Take presentation-making, for example. It's a skill we have to train our juniors into: what makes a compelling presentation. A lot of this is under the surface. Now, with AI potentially taking over these tasks, we're worried about how our juniors will develop that critical eye, that intuition for what works. It's not just about the end product, it's about the process of learning and developing judgment. We're grappling with how to ensure our future leaders still develop these crucial tacit skills in an AI-augmented environment." "I mean, to be honest with you, [...] the challenge is, do we even know exactly what our capacity is? [...] Do we actually have the baseline data in the right format to say [...] how many engineers we have? Basically, the simple question of what it would take to deliver this project in a codifiable manner, which says [...] a person in my team is an engineer of this level, he has skills in this and this and this, and he can do a project of this complexity. I'm afraid we don't have that. [...] So, I feel like there's a bit of catch-up that is needed to even get to a certain level of base data brilliance to enable this [leveraging of generative AI]. [...]" "Let's take the example in a council. When people do, let's say, cash allocation, when a voice comes and payment and the internet, but a lot of things are still done manually. And then comes AI too, which does a lot of this kind of things. When you do implement this kind of things, where do you start? You are asking the accountants of the base,

	secret tacit knowledge and they ask you, ‘Can you share your cheat sheets with us?’ Because, you know, there’s some things are obvious, rules, but there’s something where you just, they just know by experience that if this customer is paying this way it’s a little bit late and what happens? Based on this kind of tacit knowledge, it can be plugged in, and this is really a part of learning. It’s not generative AI, it’s alteration AI.” “At the zoo over there, the keepers had extraordinary tacit knowledge. So they could go close to their charges and they would know something was wrong. So the best pregnancy test for elephants at the zoo was not chemical testing or anything, it was actually a keeper coming in and just—and you could ask that keeper and you’d say, well actually she would say I don’t know but she’s pregnant.” “Soldiers call atmospherics. [...] They go on patrol and they’d always have somebody there who’d be able to walk into a village and say, something’s not right here. And they’d know. [...] How can AI do atmospheric? That’s the great question.” “I think my opinion is the special sauce, if you like, which I don’t think Gen AI will ever be me. I have played with this notion of feeding in all of my reports and studies and surveys and everything else that I’ve done for clients over the last seven or eight years into a box. So I could put [REDACTED] in a box if you like, but I haven’t done that yet. But I still don’t think it would come up with the topicality and the relevance that I bring.”
Experimentation Capacity	“Yeah, the possible reaping of benefits... China seeds clouds in order to prevent droughts, and America first used cloud seeding. They had one terrible event where they seeded clouds and it possibly intensified a hurricane and killed a lot of people, but China regularly uses that. You need a pretty good forecast, and that’s something where AI could help a lot. They seed clouds to prevent droughts in certain areas, whether that’s hydrotown areas or crop areas. So, things like that where people are willing to do something slightly outside of the norm to reap the benefits of the forecast. But something says better, can we save the sun? I think it will change the dynamics, but the dynamics are already changing anyway. I think it will accelerate those changing dynamics.” “As long as the entire system is stuck with all of these processes and procedures that they implemented years ago, so finally, there is also a big barrier that we cannot properly enter into it because as long as this [new technology] provider doesn’t [provide any backward capability], then you’re stuck with the [old] system.” “The ability to experiment freely with new technologies like AI isn’t something you can just decide to have. This has nothing to do with AI, in fact. It’s built up over years, maybe decades, of fostering a culture that embraces risk and learning from failure. By the time a disruptive technology like generative AI comes along, it’s almost too late to start building that flexibility. It’s something that happens over the course of an organization’s history.” “They are happy for people to spend time playing with AI models in order to have some confidence in their output. [...] They’re more excited to experiment. [...] A lot of companies in [REGION], there are much more pro-risk or less risk-averse using AI than I think we are in [REGION].” “Their willingness to experiment. [...] Their willingness to experiment, to accept the potential for change, to embrace change—they’re all together, aren’t they.” “The only way you’re going to learn is by trying. [...] You have to start using it, you have to start using it as a partner. You have to start asking the questions and you have to start, you know, checking the waters in a risk-free way.” “What I would be able to do to make, I think, the biggest difference is a couple of things. One was our position in the market as a challenger. [...] Second was a significant amount of resource in this area. Our digital transformation outfit, there was 800 of us. It’s not a couple of guys doing it on a Friday afternoon. It’s a big part of a big company. Third was a senior exec who really got it and loved it, who’s passionate about it. [...] So I think you do need the evangelists, but a credible evangelist who gets it and is doing something with it that can then direct activities. Otherwise, it sort of stays at the level of the coder who’s worried about his job.”
Pattern Recognition (as Vulnerability)	“Pattern recognition is a massive part of what we do. So I think that’s a big issue.” “The big fear, and it is a fear that we have as an SME, is that for the bigger and larger companies they may be able to leverage AI and they may be able to get that to do things that we couldn’t do and render us useless as such. And there is a concern about that. And I think it’s how we will use that AI to provide better information, better results that’s going to keep you ahead. But if you stay back, I think you’re dangerously getting overtaken.” “Because I think we’re still in the process of understanding GenAI in terms of it being the knowledge-bearer, if that makes sense. So, I don’t know how far we in the education sector are going to be utilizing GenAI as an organism to design products, design content, because we still need experts. Our expertise is in our brain through our experience and knowledge and qualifications. So I think it’s really about how do we manipulate GenAI to act as one of our assessment experts or markers—almost like a

	robotic system that does exactly what our expert does. And to what extent can we then control that robotic system so that it doesn't go haywire or manipulate the content in such a way that it causes some level of malpractice?"
Modularity (as Exposure)	"The more modular our processes become, the easier it is for competitors to replicate what we do. It's just the reality of our business." "We have tried to strive for a more modular structure. We're acutely aware that GenAI could potentially replace entire functional units. For instance, our data analytics module could be significantly disrupted by AI capabilities. This modularity makes us more adept, more agile, but also more vulnerable to plain imitation." "If you can break things down into very simple instructions, then it's very easy to ask gen AI for those simple, single point answers. So if it's a very convoluted story, then your trust in AI becomes less because you don't expect it to graph it, or actually, maybe you don't expect to be able to ask the question in the correct manner." "There was talk of us being acquired by a private equity firm many years ago. In response, we made intentional decisions to make our structure less modular to decrease the risk of being picked apart piece by piece—you know, to avoid becoming a body shop."
Pattern Recognition (as Scalable Advantage)	"I think the other thing is now with AI is protein folding and CRISPR. The thing about protein folding is that even though you know what the atom, what the atomic structure is in a polypeptide, it's still, people don't know why it folds in the way it does. And the reason DeepMind did it is they just piled piles of data in, it can calculate things at a monumental rate, and it can work out the following. This is critical to the rapid production of drugs, absolutely critical to the rapid production of drugs. The question is, are we going to allow, goes back to rhesus monkeys and human beings, are we actually going to get rid of some of these laborious testing stages in human beings by trusting the AI to be right? We're already facing this. The tests that came out last week were basically looking at breast cancer in women and the AI found more cases, early-stage cases of cancer than the human interventions." "Looking at a specific problem that came up and that was the problem of managing crowds at the [INTERNATIONAL EVENT] and vast numbers of people milling around and very high levels of threat of terrorism by a very small number of people. [...] The number was something like 30,000 people would be needed to look at crowds. [...] So if you take a small number of trained people, you can then magnify the effect of that small number and apply it across." "If you've got a radiologist who's looking at a scan and is trained to spot something that is slightly different, then generative AI can look at all of those pictures all over the world, massive numbers of images more than any human could possibly look at in their entire lifetime. But using the pattern recognition identified by the human to then program that, it can then throw up examples. And I think that's a brilliant example of how gen AI can be used to augment what people are doing rather than replace it." "Organizations with tacit or proprietary knowledge will be able to use these tools on their walled garden and make it even better and do something unique with it. And also young entrepreneurs will be able to access a tool beyond their reach 10–15 years ago if they've got a good idea. [...]"
Modularity (as Integration Opportunity)	"Our modular structure allows us to experiment with GenAI in specific units without risking the entire organization. We can quickly identify where it adds the most value and scale those applications across other modules. This ability to rapidly iterate and deploy gives us a significant competitive edge." "The impact and adoption will be based on how in the future we can best integrate individual modules with different AI tools. And obviously, as it develops, we're going to have different tools that are very well-fulfilling functions. And that's where really the next evolution is going to be." "So I think that sometimes, while these companies may have big treasure troves of data, the size may hinder their ability or hamper their ability to very quickly adjust and change to, as we're seeing, things changing on a daily basis for industries and for clients. So, I think that companies that are able to, even if they're very large, maybe modularize themselves in a way that, while they have a consistent pipe of consistency across the business they have, they're able to allow small parts. It's almost like operate as a network and allow small parts of the organization to go at speed and to change and deliver new products and new services and get closer to the customers at speed. So, I think companies that are able to do that will be the ones that will take a leadership position, not necessarily just the fact that they're incumbents, because I think incumbents could be hampered and held back."
Pattern Recognition (Reconfiguration)	"Yeah, I think that comes back to the thing of not wanting to miss an opportunity. In our sector, in the same way as any other industry, you're always wanting to keep up with your—in charities, we don't like to think of it like that, but keep up with your competitors. So, if they're doing certain things, then we want to embrace the same opportunities. And, you know, AI solutions are everywhere. So, yeah, I guess that's why, as well, we're just taking the approach of finding what is out there, contributing to

	different discussions to kind of stay conscious of what's happening. But, yeah, there is a slight element where we do think it's going to have this big impact, but we're not quite sure yet beyond some sometimes quite basic operational efficiencies, beyond that, what the potential for us might be." "The cost reduction is especially [in] the current economic climate, most of the big corporations are affected. And there is another element where this is going into their innovation lab is: how can we look at [a] use-case-driven approach, as you're rightly saying, to create growth. So, growth could be how to optimize market, how do I do sales better using the set of platform and tools and taking a use-case-driven approach. But currently, only based on my personal experience, the cost reduction and growth are the two different use cases that are being looked at. Yeah, so the biggest change that I have seen in this current engagement is they have been asked to reduce 10% of their IT budget. And AI has not been touched. The budget is not just reduced from other areas. The idea is, can we reduce it, but then we move it into this bucket. And this is like a very big challenge for them is about reducing costs, but also keeping ahead of times." "We will be thinking of added value, and productivity per person will increase. And we will see things which we used to do, not to do anymore [...]. Sectors will be differently affected, so we don't know how much of this AI will penetrate consulting, or legal services, or other kind of services. It could be more intrusive in other parts, and less intrusive in others. In law, you could just have AI negotiating contracts by itself, so you don't need lawyers for it, or going to court, and court being automated eventually, and not needing to be doing a lot of litigation work for that. Possibly, this could be the end of lawyers, eventually, but we don't know where this is going to end." "I think the third sector in general is more of that approach of tentative and not wanting to be left behind, but not yet doing as much, partly to do with resource, not having as much to put into AI. But yeah, certainly that tension between [a] cautious approach and wanting to try new things and follow the rest of the curve within the sector."
Modularity (Reconfiguration)	"Thinking about media and advertising and data industries, how there is a concern that the continued improvement in operational efficiency will change the dynamics of the large advertising agencies. [...] The question becomes, instead of throwing all my costs into trying to fix that infrastructure, does that make sense, or do we stop doing that, leverage it more for efficiency [...] and actually allow the disaggregated, almost small companies, several small companies, as opposed to trying to be one big company that have a competitive advantage of scale." "In our industry, the modular firms are already racing ahead with GenAI integration. We're seeing new entrants leveraging AI to unbundle traditional financial services, creating highly specialized and efficient modules. Everyone needs to act fast or be left behind." "Our processes are so interconnected that it's hard to see how GenAI could simply slot in and replace any significant part. We're excited about this technology but can't really see a clear use case. Nor can our competitors, from speaking with some of them." "While we recognize the potential of GenAI, our industry's structure and regulatory environment make rapid, transformative change unlikely. We anticipate more gradual adoption and evolution rather than outright disruption." "And I think it might be also again different for different consulting firms, but just talking about the one I know [...] We have three layers of what we would call analysts. So, we have the associate that is joining [REDACTED] and that can ultimately become an equity partner. He or she is being supported by a knowledge team. [...] And then we have specific research teams that churn [out] company reports and things like this. [...] What will be impacted is our knowledge teams and our research teams. But this has been separated, I would say, like 15 years ago where we said, "Actually, the resource junior associate is so valuable that we don't want him or her to go into company reports, to compile the two pages [on] logistics market 1995 as [REDACTED] did in his first year." [...] Because these are the guys who then need to write the slides. These are the guys who need to do their analysis. These are the people that later will become project leaders. And that's the pyramid you need to have." "We're looking at GenAI more as a tool to enhance our existing processes rather than fundamentally change how we operate. Our integrated structure makes it challenging to isolate and experiment with AI in specific areas."

C3. Further Details into Investigation

While our manuscript presents the main insights from Stage 3, here we elaborate on how we derived these findings. We find that executives' perceptions of GenAI's disruptive impact are fundamentally shaped by how they conceptualize their business's core value proposition - the essential value they provide to customers. When GenAI capabilities align closely with this core value proposition, executives perceive greater displacement threat; when they view their value proposition as orthogonal to GenAI capabilities, they perceive GenAI primarily as complementary. This value-proposition conceptualization influences perceptions of (1) displacement threat, (2) differentiation potential, and (3) expectations of overall industry disruption.

Dimension 1: Threat of Firm Displacement

Our analysis revealed that executives conceptualize their value propositions in different ways that significantly influence their perception of displacement threat from GenAI. Four primary value-proposition conceptualizations emerged from our data: regulatory intermediation, provision of output accuracy, relationship management, and technical expertise.

Regulatory intermediation

Executives who defined their core value proposition as navigating regulatory frameworks and ensuring compliance perceived a substantially lower threat of displacement. In these contexts, the value is intrinsically tied to the regulatory environment itself, creating a buffer against GenAI displacement:

A2a (banking): We [in financial services] will test everything because we are regulated. And there are a couple of bottlenecks. [...] So, one, of course, is the AI regulation itself. How would FSFCA, who regulates us, see it? They have not defined it yet. [...] I think there's not one regulation where we can talk and find things. [...] Of course, the big tech companies don't like it because they've been asked now to go basically submit everything so that we can validate. [...] They don't want to extract things and run through it because they want to push it to us.

Executives in highly regulated sectors highlighted that their core value involves the assumption of legal liability—a function that GenAI cannot replicate:

A2a (financial services): Technically, even humans as brokers, we are not allowed to advise. We can recommend or we can provide information. [...] And now, not a human, you ask an avatar who is picking up 20 tons of information from the dark web and giving advice [...] And you buy it. And you lose all your money. Who is responsible? This is a huge issue.

These executives noted that outside consultants often misunderstood the centrality of regulatory compliance to their value proposition:

A2a (financial services): I think some other consulting firms, [consulting firm], for example, they're just pushing [...] [Consulting firm] is busy here trying to get more business. Get that reporting. But they are basically saying that [we can offer] advisory by robots. Robo-advisory and stuff like that. But this falls under a regulatory grey area, maybe even a dark one.

Provision of output accuracy

Executives who conceptualized their value proposition as delivering high-precision, accurate outputs perceived their offering as inherently resistant to GenAI displacement. These executives emphasized how their business success depends on consistently producing results with minimal margin for error:

A3a (healthcare consultancy): Having an advertising campaign that's a bit off is one kind of risk. But having a bridge that is a bit off, or receiving [inaccurate] advice on your health? That is a completely different story altogether. [...] In these industries, precision is paramount and even minor mistakes can have serious consequences.

A2a (financial services): When it comes to perhaps in the reality of the stats, financial services, for example, is not picking up on it. Financials need precision and AI has a habit of being drastically wrong a lot of the time. That might be accurate if you want something between 1 and a million. But if you want something that is between 1 and 1.1—no, it's not good enough.

In sectors where output accuracy is non-negotiable, executives perceived significant barriers to GenAI adoption due to current technological limitations in ensuring consistent precision:

A2a (insurance): Do you think I'm going to buy and invest my money based on just consulting AI? [...] I would still go out and look for a person to give me good advice [...] Not because the AI is better than the humans—maybe the humans are the same—but because it's a person I trust.

Relational management

Executives who defined their core value proposition as managing and developing client relationships perceived less threat from GenAI displacement. These executives viewed their primary value as creating and maintaining interpersonal connections and trust that GenAI cannot replicate:

A2a (restructuring): I come from a very knowledge-based [...] professional service in my life, IT service most of my life. [...] Where a lot of your work is very much around, you know, knowledge, data, content. [...] So, I think in many of those industries, operational efficiency, over time as people are learning, will lead to massive initial transformations because costs will come down significantly and people will be able to use their time on more human-type relationship areas of focus.

These executives recognized that while GenAI might enhance efficiency in knowledge-oriented tasks, the human elements of building trust, managing expectations, and navigating interpersonal dynamics remained central to their value proposition and largely beyond GenAI's current capabilities.

Technical expertise

Executives who conceptualized their value proposition primarily as providing technical expertise and specialized knowledge perceived a substantially higher threat of displacement from GenAI. These executives recognized that GenAI's capabilities overlap significantly with their core value:

A1a (engineering consultancy): [...] our industry is all about the application of technical knowledge [...] I think the big fear, and it is a fear that we have as an SME, is that for the bigger and larger companies, they may be able to leverage AI [...] and render us useless as such, and there is a concern about that [...].

A1a (professional services): If your business still relies on technical knowledge, then you're toast, basically. We don't see any benefit, frankly, in trying to hang on to technical knowledge nowadays. It is no longer a differentiator and won't distinguish your business in any way, shape, or form.

Some executives in knowledge-oriented firms showed a nuanced understanding of which aspects of their value proposition were susceptible to displacement and which were not:

A3b (meteorology firm): The hierarchy of firms is going to shift in our industry for sure. Those firms with technical expertise and sophisticated digital infrastructure are going to realize their differentiation has been gone [...]. We're not going to be completely substituted, we [humans] are still needed to manage statistical chaos [...] Differentiation is going to more likely center around who best manages this chaos.

These observations highlight how executives who define their core value proposition primarily as providing technical expertise perceive a much higher threat of displacement from GenAI, prompting many to reconsider their fundamental value offering:

B1 (professional services): On the technical knowledge piece [...] there's been more and more contracting-out to specialist organizations. So, I don't think any of us probably appreciate what the balance is between in-house versus out-of-house, as it were.

Dimension 2: Leveraging GenAI for Firm Differentiation

Our analysis revealed that executives perceived certain value propositions as particularly well positioned to be enhanced by GenAI, creating opportunities for competitive differentiation. Three primary value propositions emerged as having strong potential for GenAI-based enhancement:

proprietary data-driven insights, human judgment-centered services, and innovation-focused offerings.

Proprietary data-driven insights

Executives who defined their core value proposition around delivering proprietary data-driven insights perceived significant opportunities to enhance this value through GenAI. These executives anticipated that GenAI could amplify their existing proprietary data assets to create superior insights and solutions:

B2 (law firm): Collaborations between pharma companies on developing a vaccine involve a lot of intellectual property, licenses, and terms. [...] To illustrate the point, imagine two law firms who have completed French nuclear energy projects. They have to navigate planning laws, indemnities, liability allocation, supply chain issues, government funding, and other external funding. But you can't just put that into ChatGPT and get an answer to draft a French nuclear energy contract. It's impossible, because the system isn't trained on that specific data, which likely exists only in a few cloud-based files. That data is intellectual property—it's incredibly valuable, high-end, and rare. And we have billions of such valuable data points all over the world.

Professional services executives, in particular, reported realizations that their historical client engagements represented a largely untapped reservoir that GenAI could help them leverage:

B2 (consulting firm): We've been collecting data from thousands of client engagements over decades, but we've barely scratched the surface of what insights we might unlock. With generative AI, we're beginning to see our historical data in a new light. [...] We're certain there's immense value hidden in our data that could revolutionize how we serve our clients.

B2 (professional services): If I take the world of, you know, financial diligence, which we are probably the world number one in doing, we have done more financial diligence exercises than almost anybody else put together. We know how to do that. We've got the historical data. We've got the metrics and we've got a bench strength of thousands of people who can help effectively train a model to do that.

As executives considered how GenAI might enhance their data-driven value propositions, many recognized that data cleanliness would be a critical moderating factor:

A2a (restructuring): Financial services typically will have enough data anyway. It's just a matter of whether the data is clean enough [...] that enables you to build a model that understands the different types of data and can be fine-tuned on that data.

A1a (technology firm): We have all this proprietary data that's incredibly valuable, but accessing it is a nightmare. The effort to merge and clean the data is enormous, and we're not even sure if we can do it successfully.

These concerns about data accessibility were particularly pronounced in sectors with complex legacy systems and regulatory constraints:

A2a (financial services): Data itself is a huge issue for our industry. Huge controversies. We still can't have a 360° view of our customers. Data is [spread] across thousands and hundreds of systems, especially in the insurance sector.

Human judgment-centered services

Executives who defined their value proposition around delivering superior human judgment and tacit expertise perceived opportunities to enhance this value by using GenAI to augment rather than replace human capabilities. These executives emphasized how GenAI could help their experts focus on higher-level judgment tasks:

A2b (construction): In our sector, tacit knowledge is the intricacies behind building a wall and then plastering it. Expert plasterers know exactly what they're doing, and it's quite difficult to replicate that with a machine, though we can do it with brickwork. There comes a point when perhaps we'll be able to do it with plaster as well [...] Our expectation is that we have the people with the best understanding of these tacit features and hence will be able to use generative AI better than others.

Interestingly, these same executives expressed concern about how GenAI might disrupt traditional pathways for developing the human judgment that remains central to their value proposition:

A1b (consulting firm): Take presentation-making, for example. It's a skill we have to train our juniors into what makes a compelling presentation. [...] Now, with AI potentially taking over these tasks, we're worried about how our juniors will develop that critical eye, that intuition for what works.

This observation highlights a paradox: while executives see GenAI as potentially enhancing their judgment-centered value propositions, they worry the technology may simultaneously undermine the development of the very capability they seek to enhance.

Some executives pointed to challenges in defining exactly what human judgment entails in their context:

B2 (manufacturing): The challenge is, do we even know exactly what is our capacity? [...] Basically, the simple question of what it would take to deliver this project in a codifiable manner [...] I'm afraid we don't have that. [...] So, I feel like there's a bit of a catch-up that is needed [...].

Experimentation capacity

Executives who defined their value proposition around their organization's capacity to experiment and adapt saw significant opportunities to enhance this value through GenAI. These executives emphasized how their existing orientation toward flexibility and experimentation positioned them to leverage GenAI for competitive advantage:

A3b (meteorology firm): Yeah, the possible reaping of benefits [...] China seeds clouds in order to prevent droughts [...] You need a pretty good forecast, and that's something where AI could help a lot. [...] So, things like that where people are willing to do something slightly outside of the norm to reap the benefits of the forecast. [...] I think it will accelerate those changing dynamics.

These executives viewed their experimentation capacity not merely as an organizational attribute but as a core value proposition that GenAI could significantly enhance. They anticipated that GenAI would accelerate their ability to explore novel solutions and approaches, further distinguishing them from more rigid competitors.

However, executives emphasized that this experimentation capacity could not be easily developed overnight:

A3a (manufacturing): The ability to experiment freely with new technologies like AI isn't something you can just decide to have. [...] It's built up over years, maybe decades, of fostering a culture that embraces risk and learning from failure.

Executives who defined their value proposition around experimentation were particularly conscious of how legacy structures could impede their ability to fully leverage GenAI:

B2 (technology firm): As long as the entire system is stuck with all of these processes and procedures that they implemented years ago, so finally, there is also a big barrier that we cannot properly enter into it because as long as this [new technology] provider doesn't [provide any backward compatibility], then you're stuck with the [old] system.

These insights suggest that firms with experimentation-focused value propositions may be particularly well positioned to create differentiation through GenAI, potentially widening the gap between themselves and competitors who lack this orientation.

Dimension 3: Perceptions of Overall Sectoral Disruption

This section explores how executives' perceptions of GenAI's impact vary based on how they define their core business role and value proposition to customers. We focus on two distinct business roles that emerged as particularly salient in our analysis: (1) pattern recognition as a central value-adding function, and (2) modularity in business process architecture. When executives view their core business role as primarily delivering pattern recognition value to clients, they perceive greater displacement threat from GenAI. Similarly, when they define their

organizational structure as highly modular, they anticipate more significant disruption potential from GenAI.

Pattern recognition

When executives defined their organization's core value proposition as interpretive pattern recognition—the ability to identify, interpret, and exploit recurring structures or trends in data, processes, or phenomena that provide value to customers—their perceptions of GenAI's disruptive potential were significantly heightened. Our roundtable discussions with executives who rated pattern recognition as central to their business success (A1a, n = 23) versus those who rated it as peripheral (A1b, n = 13) revealed striking differences in perceived threats and opportunities.

Executives whose business role centered on delivering pattern recognition value to clients expressed substantial concern over GenAI's potential to replace human-driven pattern-recognition functions:

A1a (pharmaceutical firm): I think the other thing is now with AI is protein folding and CRISPR [DNA modification technology]. [...] And the reason DeepMind did it is they just piled piles of data in; it can calculate things at a monumental rate, and it can work out the following. This is critical to the rapid production of drugs [...] The question is [...], are we actually going to get rid of some of these laborious testing stages in human beings by trusting the AI to be right? [...] The tests that came out last week were basically looking at breast cancer in women and the AI found more cases, early-stage cases of cancer than the human interventions.

In contrast, executives who positioned their business value outside of pattern recognition showed minimal apprehension about GenAI's disruptive potential:

A1b (charity research institute): In our sector, in the same way as any other industry, you're always wanting to [...] keep up with your competitors. [...] AI solutions are everywhere. [...] But, yeah, [...] we do think it's going to have this big impact, but we're not quite sure yet beyond some sometimes quite basic operational efficiencies [...] what the potential for us might be.

This divergence extended to how executives viewed GenAI's potential to create competitive differentiation. Those who defined their business role around pattern recognition saw GenAI as a transformative tool that could fundamentally enhance their core value proposition:

A1a (education technology firm): Our expertise is in our brain through our experience and knowledge and qualifications. So I think it's really about how do we manipulate GenAI to act as one of our assessment experts or markers, almost like a robotic system that does exactly what our expert does, and to what extent can we then control that robotic system [...].

By contrast, executives whose business roles were less centered on pattern recognition primarily framed GenAI as offering incremental operational improvements:

A1b (logistics & manufacturing): Yeah, so the biggest change that I have seen in this current engagement is they have been asked to reduce 10% of their IT budget. And AI has not been touched. [...] How can we transform the AI as such, when you talk about their team and the data scientists and the team that they have built, they are still doing smaller use cases. [...] I see this will become an important or probably the most important area for [the] C-suite to invest [in] heavily when you talk about transformation.

The implications for industry-level disruption similarly varied based on executives' perception of their business role. When pattern recognition was central to their value proposition, executives consistently articulated scenarios of substantial industry transformation driven by GenAI:

B2 (law firm): We will be thinking of added value, and productivity per person will increase [...] And sectors will be differently affected [...] In law, you could just have AI negotiating contracts by itself, so you don't need lawyers for it, or going to court, and court being automated eventually [...] Possibly, this could be the end of lawyers, eventually, but we don't know where this is going to end.

Conversely, when pattern recognition was peripheral to their value proposition, executives generally framed industry changes in more incremental terms:

A1b (charity research): I think the third sector in general is more of that approach of tentative and not wanting to be left behind. [...] But I also think there are some smaller charities who are maybe more able to take the risk and just try things, whereas, you know, we are taking a very risk-based approach to trying new things because there's so much that's unknown about AI. [...] I would say the difficulty for us is we don't have a lot of those use cases that you're talking about.

Modularity

The degree of modularity in an organization's business architecture emerged as another key dimension through which executives interpreted GenAI's potential impact. Modularity refers to the extent to which a firm's business processes, services, or products can be separated, recombined, and operated independently (Baldwin and Clark, 2000). Organizations that define their value proposition around highly modular business architectures perceived both greater displacement threat and greater differentiation potential from GenAI. To investigate this further, we conducted a roundtable with a mix of respondents who reported high organizational modularity ($n = 8$) and low modularity ($n = 7$). Executives from highly modular organizations consistently perceived greater vulnerability to GenAI-driven displacement:

A4 (manufacturing): The more modular our processes become, the easier it is for competitors to replicate what we do. It's just the reality of our business.

A4 (financial services firm): We have tried to strive for a more modular structure. We're acutely aware that GenAI could potentially replace entire functional units. For instance, our data analytics module could be significantly disrupted by AI capabilities. This modularity makes us more adept, more agile, but also more vulnerable to plain imitation.

In contrast, executives from less modular organizations, where value is delivered through tightly integrated systems, expressed less concern about GenAI's displacement potential. Simultaneously, executives who defined their business architecture as highly modular saw greater potential to leverage GenAI for competitive advantage:

A4 (software firm): Our processes are so interconnected that it's hard to see how GenAI could simply slot in and replace any significant part. We're excited about this technology but can't really see a clear use case. Nor can our competitors, from speaking with some of them.

A4 (technology firm): Our modular structure allows us to experiment with GenAI in specific units without risking the entire organization. We can quickly identify where it adds the most value and scale those applications across other modules. This ability to rapidly iterate and deploy gives us a significant competitive edge.

Executives from less modular organizations, however, primarily viewed GenAI as offering incremental improvements rather than transformative potential:

A4 (healthcare provider): We're looking at GenAI more as a tool to enhance our existing processes rather than fundamentally change how we operate. Our integrated structure makes it challenging to isolate and experiment with AI in specific areas.

As with pattern recognition, there was a marked difference in perceptions of industry-level disruption between the high- and low-modularity groups. Executives who defined their business architecture as highly modular anticipated more radical industry transformation:

A4 (financial services): In our industry, the modular firms are already racing ahead with GenAI integration. We're seeing new entrants leveraging AI to unbundle traditional financial services, creating highly specialized and efficient modules. Everyone needs to act fast or be left behind.

B2 (professional services): It's almost like operate as a network and allow small parts of the organization to go at speed and to change and deliver new products and new services and get closer to the customers at speed. So, I think companies that are able to do that will be the ones that will take a leadership position [...].

In contrast, executives from less modular organizations expected more gradual industry evolution:

A4 (energy utility): While we recognize the potential of GenAI, our industry’s structure and regulatory environment make rapid, transformative change unlikely. We anticipate more gradual adoption and evolution rather than outright disruption.

While the current section contains condensed versions of certain quotes, Table D2 provides a more comprehensive selection and uncompressed versions of participant statements. These extended quotes offer readers additional insight into the raw data that informed our analysis.

Appendix D

Stage 4 – Survey, One-to-One Interviews, Roundtable (Apr–Jul 2025)

D1. Second Round Survey

We administered a newly designed Qualtrics survey to director-level executives between April and June 2025, approximately 14 months after the Stage 2 survey. The instrument was distributed to all Stage 2 respondents for whom contact information was available, supplemented by additional IoD members recruited through the same liaison whos facilitated earlier stages. Of 63 total responses, 41 were fully completed, representing our analysis sample. Twenty-seven of the 41 respondents (65.9%) had also completed the Stage 2 survey, enabling within-person comparison across a 14-month interval.

The survey probed GenAI's competitive implications through several distinct modules: (1) the relative likelihood of GenAI enhancing versus displacing the respondent's core offerings; (2) a hypothetical automation scenario assessing expected competitive erosion if GenAI could automate 50 percent of professional tasks; (3) value proposition decomposition across six dimensions—technical expertise, relationship management, compliance and regulation, innovation, process efficiency, and creative services—capturing both current importance and expected five-year GenAI usage; (4) organizational characteristics including GenAI-specific roles, hours spent with GenAI tools, perceived task alignment, and modularity; and (5) changes in customer mix and revenue composition over the preceding year. All items used five-point Likert scales unless otherwise noted.

Enhancement versus displacement

Directors rated the likelihood that GenAI could enhance their core offerings ($M = 3.78$, $SD = 1.33$) substantially higher than the likelihood it could replace them ($M = 2.44$, $SD = 1.16$). A paired-samples t-test confirmed this difference was statistically significant ($t(40) = 6.35$, $p < .001$, $d = 1.00$). Two-thirds of directors (65.9%) rated enhancement as more likely than displacement, 31.7% rated them equally, and only one director (2.4%) rated displacement higher. The distribution of enhancement ratings skewed toward higher values, with 68.3% rating enhancement at 4 or 5, while displacement ratings concentrated at the lower end, with 58.5% rating displacement at 1 or 2. The table below reports the full distributions.

Rating	Enhancement n (%)	Displacement n (%)
1 (Very unlikely)	4 (9.8)	10 (24.4)
2	4 (9.8)	14 (34.1)
3	5 (12.2)	7 (17.1)
4	12 (29.3)	9 (22.0)
5 (Very likely)	16 (39.0)	1 (2.4)
M (SD)	3.78 (1.33)	2.44 (1.16)

Notes: N = 41. Paired-samples $t(40) = 6.35$, $p < .001$, $d = 1.00$.

Automation scenario

When asked to imagine that GenAI could automate 50 percent of all professional tasks within their sector, directors rated expected competitive erosion below the scale midpoint ($M = 2.59$, $SD = 1.02$; one-sample $t(40) = -2.59$, $p = .013$). The modal response was 3 (moderate erosion, 39.0%), with 82.9% of directors anticipating moderate erosion or less. Only 17.0% anticipated substantial or severe erosion.

Value proposition decomposition

The table below reports directors' ratings of current importance and expected five-year GenAI usage across six value proposition dimensions, along with the gap between importance and expected usage.

Dimension	Importance M (SD)	5Y GenAI Usage M (SD)	Gap	t (p)
Technical expertise	4.07 (1.06)	3.59 (1.26)	0.49	2.46 (.018)
Relationship mgmt	4.20 (0.84)	3.07 (1.27)	1.12	5.42 ($<.001$)
Compliance/regulation	3.59 (1.18)	3.56 (1.16)	0.02	0.13 (.900)
Innovation	3.56 (1.30)	3.51 (1.31)	0.05	0.27 (.789)
Process efficiency	3.61 (0.95)	4.00 (1.16)	-0.39	-1.79 (.081)
Creative services	3.17 (1.30)	3.56 (1.30)	-0.39	-2.51 (.016)
Observations				41

Notes: Paired-samples t-tests with exact p-values in parentheses. Importance: 1 = not at all important, 5 = extremely important. GenAI usage: 1 = no usage, 5 = extensive usage. Gap = importance minus expected usage; positive values indicate importance exceeds expected GenAI usage.

Relationship management exhibited the largest gap between importance and expected GenAI usage (gap = 1.12, $p < .001$): directors rated it as the most important dimension of current value creation ($M = 4.20$) yet expected it to receive the lowest GenAI usage over five years ($M = 3.07$). Process efficiency and creative services showed the reverse pattern—expected GenAI usage exceeded current importance—suggesting directors anticipated GenAI's primary contribution in dimensions they rated as moderately important rather than in the dimensions they considered most central to their value propositions.

Predictors of competitive expectations

Directors who spent more time with GenAI tools perceived greater enhancement potential ($r = 0.60$, $p < .001$) but did not perceive greater displacement risk ($r = 0.29$, $p = .064$) or greater expected erosion under the automation scenario ($r = 0.08$, $p = .633$). A similar asymmetry appeared for perceived task alignment: directors who rated GenAI as more aligned with their professional tasks perceived greater enhancement ($r = 0.56$, $p < .001$) but not greater displacement ($r = 0.19$, $p = .229$).

Organizations that had created GenAI-specific roles—29.3% of the sample (24.4% adding responsibilities to existing positions, 4.9% creating dedicated functions)—reported dramatically higher enhancement expectations ($M = 4.92$ versus 3.31 , $d = 1.72$, $p < .001$), while their displacement expectations did not differ significantly from organizations without such roles ($M = 2.67$ versus 2.35 , $d = 0.28$, $p = .463$). This pattern—direct experience and organizational investment selectively amplifying opportunity beliefs without revising threat beliefs—appeared consistently across all three predictors examined.

The table below reports bivariate OLS regression results with heteroskedasticity-robust standard errors (HC3) for the three main outcome variables.

	Enhancement	Displacement	Automation Erosion
Hours spent with GenAI	0.559 ($<.001$)	0.239 (.107)	0.055 (.677)
GenAI task alignment	0.786 ($<.001$)	0.235 (.366)	0.360 (.162)
Modularity	0.282 (.113)	0.124 (.489)	−0.019 (.899)
New roles (ordinal)	1.216 ($<.001$)	0.216 (.620)	−0.015 (.965)
Observations	41	41	41

Notes: OLS coefficients with exact p-values in parentheses; robust (HC3) standard errors. Each cell reports a separate bivariate regression.

In a multiple regression including hours spent, task alignment, and modularity simultaneously, hours spent remained the strongest predictor of enhancement expectations ($b = 0.38$, $p = .023$), with task alignment marginally significant ($b = 0.47$, $p = .086$) and modularity non-significant ($b =$

0.05, $p = .794$). The combined model explained 43.6% of the variance in enhancement expectations ($R^2 = .436$). For displacement, no predictor reached significance in the multiple regression ($R^2 = .088$), and for the automation scenario, only task alignment approached significance ($b = 0.45$, $p = .076$; $R^2 = .131$). The contrast between the well-specified enhancement model and the poorly specified displacement model further illustrates the asymmetric relationship between GenAI experience and competitive expectations.

Cross-wave panel analysis

Twenty-seven directors completed both the Stage 2 (March 2024) and Stage 4 (May–June 2025) surveys, enabling within-person comparison over a 14-month interval. Matched and unmatched directors did not differ on any Stage 4 outcome variable: enhancement ($M = 3.85$ vs. 3.64 , $p = .671$), displacement ($M = 2.37$ vs. 2.57 , $p = .624$), or automation erosion ($M = 2.59$ vs. 2.57 , $p = .949$).

The matched panel replicated the full-sample enhancement–displacement asymmetry ($M = 3.85$ vs. 2.37 ; paired $t(26) = 5.74$, $p < .001$, $d = 1.11$). Cross-wave correlations between Stage 2 competitive expectations and Stage 4 outcomes were generally weak: Stage 2 industry importance, distinctiveness potential, and financial viability expectations showed no significant associations with any Stage 4 outcome. The one exception was Stage 2 disintermediation likelihood, which predicted Stage 4 enhancement expectations ($r = 0.38$, $p = .049$).

Stage 2 firm characteristics, by contrast, showed selective cross-wave associations. Ratings of technical knowledge importance in Stage 2 predicted enhancement expectations 14 months later ($r = 0.54$, $p = .004$), while ratings of relationship knowledge importance were associated with lower displacement expectations ($r = -0.34$, $p = .079$) and lower expected automation erosion ($r = -0.35$, $p = .071$). The table below reports the full correlation matrix.

S2 Characteristic	S4: Enhancement	S4: Displacement	S4 Automation: Erosion
Technical knowledge	+0.54 (.004)	−0.06 (.774)	−0.01 (.975)
Pattern recognition	+0.04 (.827)	+0.08 (.697)	−0.02 (.930)
Relationship knowledge	+0.07 (.721)	−0.34 (.079)	−0.35 (.071)
Market insight	+0.26 (.188)	+0.10 (.612)	+0.23 (.250)
Tacit knowledge	+0.21 (.290)	−0.26 (.184)	−0.16 (.425)
Certification	−0.17 (.398)	+0.05 (.813)	−0.06 (.751)
Referrals	−0.08 (.693)	−0.20 (.318)	−0.09 (.662)
Observations	27	27	27

Notes: Pearson correlations with exact p-values in parentheses. Stage 2 ratings measured March 2024; Stage 4 ratings measured May–June 2025.

These cross-wave patterns suggest that the frames through which directors filtered GenAI's competitive meaning—anchored in how they characterized their firms' value propositions—proved more stable than any specific competitive expectation. Directors whose firms relied on technical knowledge saw greater enhancement potential over time; directors whose firms relied on relationship knowledge saw lower displacement risk. The weak cross-wave correlations between Stage 2 and Stage 4 competitive expectations, combined with the stronger associations between Stage 2 firm characteristics and Stage 4 expectations, are consistent with interpretive frames that shape how directors process new information about GenAI rather than with simple test–retest stability of expectations.

D2. One-on-One Interviews

While our follow-up survey in June 2025 provided encouraging signs of perceptual stability—most notably the consistency in how executives ranked their firms' competitive positioning relative to our Stage 3 findings—we sought to probe more deeply into whether, and how, executives' views on GenAI competitive implications had evolved over time. Survey responses, while useful for identifying surface-level trends, could not fully capture the richness of directors' reasoning or reveal how they reconciled increasing GenAI capabilities with ongoing confidence in their firms' defensibility.

To address this, we conducted a series of 45-minute one-on-one interviews with executives who had indicated a willingness to be contacted for further participation. These interviews aimed to go beyond the quantitative instrument, offering an opportunity to explore apparent contradictions, deepen our understanding of mental models, and assess the durability of executives' narratives in the face of technological change. This qualitative stage allowed us to triangulate with earlier data, strengthen confidence in the interpretation of survey results, and explore whether any latent shifts had begun to take root. The sections that follow outline our interview protocol, participant demographics, and the main themes that emerged.

Between May and July 2025, we conducted semi-structured interviews with 22 directors, of whom all were recorded and transcribed for analysis. Interviews lasted between 30 and 60 minutes and followed a protocol designed to elicit (1) directors' updated assessments of GenAI's capabilities and trajectory, (2) whether and how their competitive expectations had changed since Stage 2, and (3) the reasoning through which they reconciled acknowledged capability expansion with their competitive assessments. We deliberately chose individual interviews rather than group settings to avoid the peer dynamics that might have reinforced earlier positions in roundtable formats. We organize our findings by theme.

Capability acknowledgment and trajectory

Directors uniformly acknowledged that GenAI capabilities had advanced substantially over the preceding year. A director at a strategy consultancy described the shift as moving from "a

sophisticated autocomplete" to something that "can genuinely reason through complex problems," adding that "the trajectory has been steeper than I expected." A director in construction and architecture reported using GenAI "quite a bit as a replacement for Google search to do market research," noting that "I don't have to pour through all the reports—I can just get the AI to do it for me and come up with a summary." An executive coaching consultant described moving from no subscription to a paid ChatGPT account within the year, observing that "it is a good productivity tool once you've got a workflow right," estimating it saved "several hours per conversation" for client summaries. A director at a B2B marketing agency described growing efficiency: "I'm getting used to how best to train it and brief it so that it gives me the best output. That's a learning curve, progressively getting better and better."

Stability of competitive expectations

Despite acknowledging capability improvements, directors overwhelmingly described stable competitive expectations. The strategy consultancy director continued: "But does that change whether clients need us? Not really. If anything, I'm more confident than I was. We've had 18 months to see how clients actually use this, and the pattern is clear—they use it for the analytical grunt work, just as we do, but they still need someone to tell them what it means and help them act on it." A director in construction and architecture stated that "nothing's made me nervous" about displacement, characterizing everything he had seen as something that "could help me be a bit more efficient, a bit more effective and productive," but that "the way we're working right now is quite human intensive and quite personality intensive, so I don't see it happening anytime soon." A director at an EdTech and consultancy firm posed a pointed question: "Has it massively enhanced the profit of the company? No. Has it massively contributed to a bigger bottom line? Slightly." He added that "there hasn't been this wholehearted ripping out and job replacement that some predicted. It is really much more gradual than that."

Workforce implications and historical parallels

Directors recognized that GenAI was reducing headcount at junior levels but consistently framed this as continuous with prior waves of automation rather than as a novel competitive disruption. A director at an accounting firm described how successive waves of technology had transformed his team composition: "When I first worked for my first practice, I was one of a team of six juniors, a semi-senior, a manager and a partner. Now we have one equivalent junior, one senior and one partner. So the work of six has become the work of one." He characterized GenAI as simply the latest iteration: "All the automation has happened at the lower end. It's the thinking part, the judgment, that's popped out at the top." A director at a venture capital firm invoked explicit historical analogies: "This is the age old argument, right? Do we really need that many railroad workers today? Probably not. And do we need any people that are able to operate horse drawn carriages? Probably not, because it's just where shifts have to happen." A director at a B2B marketing agency identified the emerging pipeline problem: "Where are we going to get our entry-level people from? Because actually AI can do most of what they can do, like the juniors in the industry, and how are they going to become seniors if they never become juniors? And this has

already happened in accountancy, management consultancy, and law." He nonetheless normalized this within a broader pattern: "We know that through every major automation exercise, jobs get destroyed, but jobs get created."

A director at a consultancy and EdTech firm described a similar pattern at the sector level: "The big four consultancies, the big five or six accountancies, they are betting heavily on AI, meaning they need less grads to come in." Yet he immediately raised the structural question this created: "I'm not sure how they're going to get their middle management. They don't bring in youngsters and young talent at the bottom. How are they going to grow them up?" This concern about the talent pipeline—shared across several interviews—treated workforce reduction as a genuine organizational challenge rather than as evidence of competitive displacement.

Reconciliation through proprietary assets and relational capital

Directors who had emphasized proprietary data or tacit knowledge as sources of differentiation described capability improvements as amplifying existing advantages. A director at a B2B marketing agency described a client—a global engineering firm—that had "basically brain-dumped the 80 years of their existence into an AI-driven database and a powerful tool," enabling instant retrieval across hundreds of complex construction projects. "You still can't do it without an engineer," he noted. "An engineer still needs to design the bridge. But a lot of the supporting information at the touch of a button has just appeared." A director in venture capital described training "a GPT separately to essentially evaluate companies the same way that I would," which he had licensed to other funds. He characterized the competitive moat as "the individual organization's core functionality and core offerings and domain expertise"—assets that GenAI could leverage but not replicate independently. An independent advisory director observed that "those very specialist boutiques are already in the business of training AI to help them make permanent the advantage that they have," framing GenAI as a mechanism for codifying rather than displacing proprietary knowledge.

Directors who had emphasized relational or implementation dimensions described capability expansion as irrelevant to the dimensions that mattered for competitive positioning. A director in construction and architecture observed that despite technological advances, "a lot of that is still trust-based, person-to-person relationship building," adding that he did not see this being displaced "anytime soon." A director at a life sciences consultancy stated that clients ultimately chose his firm because of "a massive element of personal relationship. We got this business because A knew B, B knew C. A trusted B, B trusted C. You can't underestimate that value. You can't probably generate that personal relationship element." An executive coaching consultant described his value proposition in terms that were orthogonal to GenAI's capabilities: "You can identify all the tasks, but the ability to choose which task to do in what order and then how to present the results in a human to human interface at the executive level is not something that AI will easily do." A director at a consultancy and EdTech firm described a similar division between visible and invisible AI use: "The front of house of consultancy will remain similar. You do have to have that credible voice.

They might, behind the scenes, do an awful lot of work using AI, but I think it's going to be behind the scenes, not front of house."

Liability, regulation, and institutional constraints

A corporate law partner captured a distinct line of reasoning that appeared across several interviews: "The technology is better, obviously. But the fundamental question hasn't changed—can it be liable? Can it carry professional indemnity? Can clients sue it when it's wrong? Those questions aren't about capability." A director at a financial consultancy in a regulated offshore jurisdiction described how regulatory requirements constrained adoption even where productivity gains were substantial, observing that "until regulators can become confident that a technical drawing done by an AI is just as good as one done by a licensed architect, it's going to be a while before they're willing to accept that substitute." An executive coaching consultant noted that his professional ethics code required data anonymization before using AI tools, adding work that partially offset productivity gains: "If they were able to guarantee that the data was much more strongly encrypted, I would be an even more avid frequent user of it." These observations suggest that institutional and regulatory constraints shape the competitive implications of GenAI independently of its technical capabilities.

Sector-level concerns persisting despite firm-level adaptation

A subset of directors maintained concerns at the sector level despite describing successful firm-level adaptations. An independent advisory director argued that the protection afforded by boutique expertise "is not a guarantee," warning that "if I invest in generative AI and I train it how to analyze my data and publicly available data, I don't see where you fit and I don't see why I should pay a premium." She predicted that knowledge work firms would need to evolve beyond their current value propositions or face insourcing by clients. A venture capital director, after describing substantial productivity gains from GenAI tools, reflected: "I've been among the rightfully or wrongfully confident groups of people who says, I can't be displaced by AI because there's too much that I bring to the table. I don't know that that's true anymore." These directors updated their views on what GenAI could do without correspondingly updating their assessment of what it meant for their own firms—instead locating their concerns at the industry or ecosystem level.

D3. External Roundtable

In June 2025, we convened an external roundtable at the collaborating institution, bringing together 27 participants including venture capitalists, technology executives, consultants, HR leaders, insurance executives, media editors, public sector leaders, and applied AI specialists. The session was designed to triangulate our interview and survey findings against a broader set of perspectives, particularly from participants who work with GenAI deployment across multiple organizations and sectors. Both sessions were conducted under Chatham House Rules to support open dialogue, and participants consented to the disclosure of their identities to provide transparency regarding the composition of each cohort.

What distinguished this out-of-sample group was the exceptional seniority of participants, with many occupying C-suite or board-level positions that afforded them unique vantage points across multiple sectors. Their discussions were notably reflective, often incorporating buy-side perspectives and cross-sector evaluations that extended well beyond their own organizations' experiences. This multi-stakeholder perspective enriched our validation, as executives could compare and contrast GenAI's impact across their portfolios, client bases, and competitive landscapes.

The out-of-sample group exhibited notable diversity across organizational types and sectors. Participants included executives from global consulting firms, financial services companies (including banking and insurance), technology ventures backed by venture capital, healthcare organizations focused on vaccine development, news and media organizations, retail and consumer goods companies, logistics firms, and non-profit research institutions.

Notably, several participants occupied boundary-spanning roles that offered unique perspectives on GenAI adoption across different organizational contexts. For instance, one participant led AI initiatives at a major insurance company, while another directed digital transformation at a global vaccine development organization. Multiple venture capitalists provided insights into how AI-native companies were approaching market entry, while executives from traditional corporations offered perspectives on incumbent adaptation strategies. We summarize the key themes. The roundtable composition is reported below.

Firm	Role / Position	Organization Type
BCG	Principal; Ambassador	Management Consulting
BCG	Principal; Ambassador	Management Consulting
BCG X AI Science Institute	Director; VP, AI and Data Science	Management Consulting (AI)
Google DeepMind	VP, Research	Technology (AI Research)
Fujitsu	Head of Western European Delivery	IT Services
Fujitsu	Senior Director, AI Research (Europe)	IT Services
Reuters	Global Managing Editor, AI Newsroom	Media
Ocado	Chief Data Officer	Online Retail and Tech
EQT Group	Chief Digital Officer, AI and Transformation	Private Equity
BC Partners	Managing Director, AI	Private Equity
Point72 Ventures	Advisor	Venture Capital
Acadian Ventures	General Partner	Venture Capital
IG Group	Global Head of Data and AI Transformation	Financial Services
KBC	General Manager, People Analytics	Banking, Financial Services
AIG	Head of AI	Insurance
Gartner	Research VP; Senior Research Director	Industry Research
Russell Reynolds	Head of Strategy and Excellence	Executive Search
Imperial Brands	Head of R&D and Innovation	Consumer Goods
Bpost	Director of HR Strategy	Postal and Logistics

Firm	Role / Position	Organization Type
CEPI	Director of IT and Digital Transformation	Global Health Organization
Faculty AI	Director, Retail and Consumer Sector	AI Consultancy
Aegis Labs	Chief Executive Officer	Technology Lab
IoD	Head of Innovation and Technology Policy	Policy Institute
Mantu	CFO	Management Consulting
Evolution Ltd	Engagement Manager	Leadership Advisory
Business School	Professor; Lead Advisor	Academia
Business School	Researcher	Academia

Note: n=27. The original invite list consisted of 30 individuals. Three dropped out on the day due to unforeseen circumstances. This workshop was conducted under Chatham House Rules. Participants consented to the disclosure of their identities for the purpose of providing transparency regarding the composition of the cohort. Business schools anonymized for peer review.

Task-level gains versus organizational returns

Participants reinforced the disconnect between demonstrated task-level productivity and measurable organizational impact. A head of AI at a financial services firm described running "more than 130 experiments" since ChatGPT's launch, with "a lot of them in production now," but added that "building the specific product or a specific task is reasonably straightforward if you have a good team, but actually ripping the full ROI from it is not as straightforward." A venture capital general partner was more direct: "I get plenty of founders telling me how awesome this stuff is going to be, but I'm looking to find examples of people that are actually doing something in production, however modest, that's actually having a genuine business impact. The more I look, the harder I'm finding that to find." A chief data officer at a technology firm captured the prevailing uncertainty: "Is it real? Everyone's asking themselves, is this real? Do I need to react or not?"

Heterogeneity across organizational types

Participants strongly endorsed the observation that GenAI's impact varies fundamentally by organizational type. A venture capital general partner argued that startups possess structural advantages: "If I'm building a company now, I'm building with AI natively. I'm not building with, oh, I have to do a change management exercise. The reasonably well funded startup probably today has more access to AI-relevant capability than a large corporate would." A head of AI practice at an applied AI consultancy identified only three organizational types generating substantial value from GenAI deployment: "well-funded scale-ups, government, and large corporates—because the cost of change is so great that they're the only groups that can afford it." A chief data officer offered a complementary lens based on prior disruption experience: "I've had quite a few conversations with media companies who are much more nervous than other industries, and I think that's because they've seen social media ruin their business model before." This observation suggests that organizational responses to GenAI may be shaped by firms' histories of technological disruption rather than by GenAI's technical characteristics alone.

Workforce implications and the talent pipeline

Concerns about workforce implications surfaced prominently, echoing the interview findings. A global managing editor at a major news organization identified workforce composition as the central challenge: "The challenge with us is not so much around the deployment of the tools. It's what our workforce is going to look like even 18 months out—the people we are going to be able to retrain, the people we're not going to be able to retrain. Frankly, that's by far the hardest bit." A venture capital partner noted that law firms "are not hiring associates anymore, because what was the job of associates—doing research? They don't need anyone to do research." A consultant challenged the conventional expectation that juniors are most at risk, arguing that "there are more and more opinions that it will replace a lot of middle management, because the role of the project leader is to scale, and a big part of middle management is frankly to be this middle level between top management and other employees." A head of AI at a financial services firm introduced a generational counterpoint, relaying her daughter's perspective upon joining a major consultancy: "I am an AI native. I'm not worried about my career. The careers I'm worried about are the middle management people who don't actually have the native AI skills."

Adoption barriers and shadow AI

Participants described a tension between the desire to control GenAI deployment and the reality that employees were already using consumer tools with sensitive data. A director at an executive search firm described fighting "quite an intense battle" to get enterprise AI solutions approved, noting that "the biggest concern people have now is not shadow IT, but shadow AI." A consultant at a professional services firm described the consequences of restrictive policies: firms "tried to build in-house tools, and because of so many safety valves and restrictions, those tools became unusable. And ChatGPT is banned because of sensitive data. But in fact, what happens is everybody is using ChatGPT from their phones, taking photos of slides and documents and processing them through ChatGPT, and it's a mess." A global managing editor at a news organization described a more successful approach: building AI tools that embedded directly into existing platforms, where "adoption goes from zero to 100 overnight because you were doing this yesterday and today there's an additional feature." An IoD governance expert reported that among directors surveyed, "70 percent are using AI daily, 30 percent are not even using it, and in terms of adoption, 50 percent haven't even adopted it—so there's significant resistance."

Appendix E

Stage 5 – One-to-One Interviews (Jan–Mar 2026)

E1. One-to-One Interviews

Between January and March 2026, we conducted semi-structured one-to-one interviews with directors who had participated in earlier stages of our investigation. Where Stages 3 and 4 had combined roundtable discussions with surveys and follow-up interviews, Stage 5 relied exclusively on one-to-one interviews as the mode of inquiry. This design choice reflected two considerations. First, by this point in the investigation—approximately two years after the initial roundtables—we sought to elicit detailed, reflective accounts of how individual directors’ perceptions of GenAI’s competitive implications had evolved over time, a task better suited to extended private conversation than to group discussion. Second, the one-to-one format allowed us to revisit each director’s specific earlier responses and probe directly for continuity or change in their assessments.

Sector	N	Percentage
Consulting	14	34.1
Financial Services	7	17.1
Technology	6	14.6
Policy	3	7.3
Construction	2	4.9
Pharmaceutical Advisory	2	4.9
Cybersecurity	1	2.4
Education	1	2.4
Government Advisory	1	2.4
Private Equity	1	2.4
Real Estate	1	2.4
Non-Profit	1	2.4
Hardware	1	2.4
Total	41	100.0

We contacted all directors who had indicated willingness to participate in further stages of the research, reaching out via email to schedule individual interviews. The invitation explained that the research team was conducting a follow-up consultation to understand how directors’ experiences with and expectations regarding GenAI had developed since our earlier interactions. Forty-one directors agreed to participate. Interviews lasted between 30 and 60 minutes and were conducted via video call, with all sessions recorded and transcribed for analysis. The sample spanned a broad range of sectors including consulting, financial services, technology, construction, pharmaceutical advisory, cybersecurity, education, policy, government advisory, private equity, real estate, non-profit, and hardware. The interview protocol was designed to elicit: (1) directors’ updated assessments of GenAI’s task-level capabilities and their personal use of the technology; (2) whether and how their expectations regarding GenAI’s competitive implications—at the firm, sector, and industry levels—had changed since their earlier participation; and (3) directors’

explanations for any discrepancy between the technology's perceived potential and its observed effects to date. For directors who had participated in multiple earlier stages, the interviewer referenced their specific prior responses to facilitate direct comparison.

Our analysis of the interview data followed an approach consistent with the earlier stages. Following transcription, we employed open coding to identify first-order concepts from directors' own language, then grouped these into higher-order themes through axial coding. Where possible, we compared directors' Stage 5 responses directly with their earlier statements from Stages 2, 3, or 4 to assess within-person continuity or change in competitive expectations over time. The sections below report the principal themes that emerged from this analysis.

Task-Level Productivity: Continued and Deepening Gains

Across the 41 interviews, directors reported that GenAI's task-level capabilities had continued to improve and that their use of the technology had deepened since our earlier interactions. Directors described a trajectory from initial experimentation to habitual reliance, with the technology now embedded in daily workflows for tasks including research, document summarization, report drafting, data analysis, and communication. A pharmaceutical advisory director described how his firm had approached the question systematically: "We have a workstream. The workstream's aim is that we should be able to achieve 10% year-on-year improvements in productivity through using AI tools, and that seems to us a fairly gettable target, because we've broken down projects into the different steps and thought, well, at a briefing stage, where could we use it? At this initial drafting stage, where could we use it? And so on. And there are some parts where you can't use it, but there's a lot of elements of tasks that we could accelerate using it." A consulting director described a concrete example of how the technology had enhanced an employee's effectiveness in ways that extended beyond simple task acceleration: "He starts to use Copilot from Microsoft. He actually presents in a different way now. He learned from the creation of the slides and creating a theme. And he's learned to tell stories much better, whereas before he would just present data... Whereas a year ago, they wouldn't have put him in the room because he would have sent them all to sleep. Real productivity gains there."

Directors with technical backgrounds offered particularly precise accounts of where GenAI added value and where it did not. A consulting director who advises on mergers and acquisitions across technology and telecommunications distinguished between the technology's capacity to handle explicit versus tacit knowledge: "People have these gigantic pools of data and expertise, but most of that data, it's like an iceberg. A good 70% of it is in the people, not actually in the databases." He noted that current GenAI tools could accelerate work within the explicit domain but could not yet access the implicit knowledge that underpins senior-level judgment. A government advisory director described the gains in IT security operations with specificity: "AI tools that come out from Microsoft, for example, are absolutely brilliant at scanning all our endpoints and spotting patterns and raising concerns... it can do that in two minutes, which would take two days minimum for an individual to be able to do." A non-profit director who had deployed the technology organization-wide offered a characteristic assessment of both its power and its limits: "It is exceptionally good

for some things, incredibly bad for others, because it is ingratiating. It will tell you what is there, but it won't tell you what's not there." He elaborated: "For plain vanilla documents, which are predominantly prose and text, you can do it. But if you throw pictures in, if you throw diagrams in, if you've got diagrams which are text-heavy, sometimes it still misses that stuff."

Several directors reported that the technology had enabled them to absorb tasks previously delegated to junior staff or outsourced to specialists. A technology and financial services director described how GenAI had compressed the prototyping cycle: "You can get an MVP to market in a few hours, whereas it would have taken a few weeks. It's almost like a new set of coloring pens—whereas once you'd have had a static slide deck just to introduce a concept, now you can build something. It's not production-ready, but now you can let somebody play with what your idea is in a much more holistic view." A construction industry consultant described how GenAI had altered the economics of his practice: "In the past, you would have outsourced that part... But now I find I can do that myself using the platform and double-checking. It meant we didn't have to hire a consultant, for example." Yet directors consistently emphasized that these gains remained task-level in character. As one consulting director put it: "That very first bit—okay, what might a report look like, just give me some fresh ideas, I've not got any blind spots—that's about 10% of the work. The other 90% is applying the expertise and the frameworks."

Stability of Strategic and Competitive Expectations

Despite the acknowledged improvements in GenAI's capabilities, directors described their strategic and competitive expectations as largely unchanged from their earlier assessments. When asked directly whether their views on GenAI's implications for their firms and sectors had shifted, the modal response was explicit stability. A policy director stated: "I think it's not changed significantly since I took part in the meetings, as you say, a couple of years ago." A financial services director was equally direct: "My view has not changed about the fact that this will not do away with experienced consultants and specialists, accountants, lawyers, et cetera, because you still have to watch and make sure that the sentiment captured is appropriate and that there isn't any bias that's crept in." A technology director working in AI research and development offered a particularly direct formulation: "For me, from what I see, nothing has changed. And in terms of the hyper-competitive—the kind of 'AI is everywhere'—it hasn't changed for me since 2024." Where the interviewer was able to reference directors' specific prior statements from Stages 2, 3, or 4, the comparisons consistently revealed continuity rather than revision.

This stability was observed across sectors and across the full range of competitive expectations. Directors who had expressed optimism about GenAI's potential to enhance their competitive positioning in earlier stages continued to hold that view, while those who had been skeptical remained so. A consulting director observed that while personal productivity had undeniably improved, the competitive landscape appeared fundamentally unchanged: "We're not seeing the need for what we do diminish... We could be more efficient in doing some of the person-draw, like the lengthy legacy site work, the data and prep type work. It's making us more efficient potentially, but so is everybody else that we're in the marketplace with." A government advisory director

similarly noted: “At the firm level, we’re not seeing any drastic change, because we still, on the whole, need the people.” Even a director whose cybersecurity firm had achieved substantial strategic gains through GenAI—reducing headcount in one division by two-thirds and creating an entirely new product line—observed that competitors had been unexpectedly slow to follow: “I’m always surprised that these bigger companies, you know, bigger than us, aren’t so forward with the AI. But the only thing I could attribute to that, without knowing anything better, is that it’s bigger, and it’s harder to get these things moving versus us being smaller and a bit more nimble.”

Where expectations had shifted at all, the direction was toward greater caution rather than heightened concern about disruption. A financial services director with a GenAI startup described an evolution in his personal understanding of the technology’s potential—“I stand at a juncture where I truly believe that autonomous execution, which is accurate, doesn’t hallucinate and is not biased and all the other kind of risks, is possible. I believe that now. Two years ago, I wasn’t close enough to be able to pass a judgment”—yet simultaneously observed that organizational adoption had barely moved: “The membership that we have, which I would put into the kind of SME bracket, is still quite—I wouldn’t say scared—I would say cautious of adoption. And that hasn’t really materially moved forward, even in the last year.” A management consultant specializing in financial transformation noted a subtle but important perceptual shift: “I think two years ago, it was a little bit less obvious. But after two years, I think there’s much deeper understanding of the difference—how it applies.” This deeper understanding, however, had not translated into revised competitive expectations; rather, it appeared to have tempered the initial sense of urgency.

Surprise at the Absence of Organizational and Sectoral Effects

A recurring theme across the interviews was directors’ expressed surprise that GenAI had not, to date, produced the organizational and sectoral effects that many had anticipated. This surprise was notable because it coexisted with directors’ undiminished conviction that such effects would eventually materialize. A financial services director captured this tension directly: “That is counterintuitive in as much as we would have expected this to have happened already, right? But by some of the more agile, smaller players—we haven’t seen it yet.” He continued: “To your key question—are we seeing evidence of a true transformation? I would say up to now, no.” A consulting director who advises on mergers and acquisitions across technology, media, and telecommunications described the gap between expectation and observation in direct terms: “At the moment, there’s no correlation between AI and employment that I can see from my chair, either as a director or as a consultant or involved in M&A or involved in just conversations around the valley. There’s like nothing.” He elaborated: “People are getting rid of jobs, but it’s not because of AI. That’s just—they’re using it as a cover. So the Silicon Valley mob are basically programmatically going through and removing management layers because they should have done it over the last five years and haven’t done it. So they’re right-sizing and they’re pointing to AI so they don’t get into trouble from a perception point of view.”

Directors in financial services offered a concrete account of the gap between technological promise and implementation reality. A director at a financial technology firm brought up the recent stream of

reports on the lack of A non-profit director who had deployed GenAI across his organization described reaching an unexpected ceiling: “We don’t know how to take it to the next level,” adding that he was sending his chief technology officer to an external program “specifically to go and learn what others have done to apply AI to organizations like us, because I don’t think we’re doing enough.” A pharmaceutical advisory director observed that even measuring the productivity gains that had supposedly been achieved proved elusive: “The extent to which we’ve seen productivity benefits is very hard to measure, we’ve found. We’ve asked Microsoft for data on our Copilot usage, but the data they provide, at least at the moment, is quite high level... But what we’re not really sure of at the moment is when we use particular tools, is it helping us get through projects more quickly? Is it enabling us overall to do more? So we don’t know about that.”

The expectation that proprietary data would serve as a decisive competitive advantage—a theme that had featured prominently in Stage 3 discussions—was revisited with notably deflated assessments. When the interviewer reminded a cybersecurity director that two years earlier “every single director in a professional service firm was like, we have all this proprietary data, we’re going to supercharge it and our competitors are toast,” she responded: “The data’s still there. It’s not as quick as people think.” An accounting director described a major software firm’s proprietary data claims as primarily “marketing.” A consulting director observed that the expected proprietary advantage had been eroded by the technology’s democratization: “A couple of years ago, it was felt like it was going to be something that could be contained and exploited by the big few. Because of the way it then deployed across—through ChatGPT, Copilot, you name it—that whole gig economy went, well, you don’t need to be a big firm to do this... So I think the access to it probably diminished their value proposition.” A consulting director who works with supply chains noted that the prerequisite of data quality had proved far more onerous than anticipated: “Two years ago, firms were spending about a month looking at clean data and cleaning up the data. And then over the last two years, I found out now that that’s gone up from one month to between three and really in the last six months, they’re spending between six and nine months making sure that they’ve got clean data. Without that, they’re not getting any of the advances that they need in artificial intelligence.”

Holding Strategic Frames Constant: Searching for Reasons Elsewhere

A notable pattern to emerge from the Stage 5 interviews was that directors who acknowledged the gap between GenAI’s anticipated and observed effects did not respond by revising their strategic expectations downward. Instead, they maintained their earlier assessments of the technology’s eventual competitive significance while actively searching for explanations as to why those effects had yet to materialize. The explanations directors offered were varied and often specific to their sectors, but they shared a common structure: the technology’s potential remained intact; the barriers lay elsewhere.

The most frequently cited explanation was organizational inertia and board-level resistance. A board director in finance described the challenge in terms of boardroom composition: “Boards, in my experience, are struggling. And I’ll just tell you, in the boards that I sit on, we have the full

range of opinions. We have people who are all in and want AI subcommittees... and then we have other people who think it's really not the board's role to worry about that because the execs are doing that." He added: "There's a certain amount of stasis in some of these boards—it's the way we've always done it." A financial services director offered a complementary account, arguing that risk aversion at the board level had created a bottleneck: "I am now seeing the absence of boards having the complete understanding of, let's call them, the challenges of AI, which you might transcribe in a different way as the risks of adopting AI. Boards are still coming to terms with what that all means." He attributed organizational caution to an incomplete understanding of the relevant risks: "The people are still, let's say, a little bit fearful of the two major risks that they perceive—one is bias, the other may be disingenuous. I don't think people have really thought through beyond those two risks to much higher-level risks." A director in private equity pointed to a structural dynamic within large firms: "You have incumbent CTOs who want to safeguard their jobs. And so they'll go through the planning cycle, the budgetary cycle, present to boards two or three years later. Maybe they'll start at the enterprise level before anything can be done... it could be self-preservation, but also you can't just change things overnight."

A second prominent explanation centered on the gap between individual adoption and organizational transformation. Directors consistently observed that GenAI use remained bottom-up—individuals integrating the technology into personal workflows—without corresponding top-down restructuring of organizational processes. An independent advisor framed this as a leadership deficit: "The fact that it's still true that it's seen as something that your tech people, your CTO, knows about and will look after for you and will take care of, means it hasn't quite made it into the engine room of most large organizations. And therefore, perhaps, hasn't quite been embraced right at the heart of where people are coming up with new ideas." She elaborated: "I don't think you've got the people who are well enough trained and well enough, familiar enough, senior enough levels in these organizations to be able to make that happen yet. Because I think the people who really play with it are quite young, and quite relatively junior, in organizations which are traditionally hierarchical." A technology and consulting director offered a structural account of why bottom-up adoption alone was insufficient: "If you just build on tasks, on existing execution systems, with people still driving everything, you cannot successfully utilize AI or agents. You cannot. It's impossible. Because we are still a limiting factor—humans." A consulting director concurred: "They've not really been embedded as part of operating models. People are playing with them, experimenting with them, but they're not part of the operating model, not part of the standard operating procedures. That's not happened yet."

A third set of explanations pointed to infrastructure and technical barriers that had proved more formidable than expected. Directors in financial services were particularly emphatic on this point. A financial technology director argued that the problem was not in the AI itself but in the organizational substrate onto which it was being applied: "People are trying to apply these complex AI engines over legacy implementations. And those legacy implementations are typically fragmented implementations running over heterogeneous messaging systems using batch-process data, which is being presented in disparate formats to an AI engine which has to consume

consistently formatted, high-quality data on a continuously streamed basis.” He described the underlying challenge as one of accumulated technological debt: “There’s an overhang here of technology debt... generations of systems that are being patched, repaired, added to, amalgamated with, and consolidated with other systems that aren’t naturally compatible... thousands and thousands of undocumented patches, which has led to this fragmented, disparate IT technology debt, which now has to be swept away if you’re going to implement the efficiencies that AI offers.” A real estate consultant who works on large-scale projects in the Middle East described a similar barrier: “A lot of businesses in the UK are mature businesses. So you have to retrospectively put a system in, which does upset the apple cart, because you’re still trying to produce your service or products, but at the same time you’re trying to modernize the company.” A private equity director pointed to the downstream consequence of integration failures: “Until you change the way the platform actually works... yes, the chatbot was developed with AI, but until the bank’s system knows exactly what the chatbot is supposed to engage with, all their background and inherent structures, it will never be effective. So it’s not the fault of the chatbot. It’s just because everything else behind it doesn’t work.”

A fourth explanation, offered primarily by directors in consulting and technology, framed the current gap as consistent with historical technology adoption cycles and therefore unremarkable. A consulting director argued: “It’s going to take five, seven years for this to ripple through, just like all technology transitions. There’s a classic curve here and we’re right at the beginning of it.” He attributed inflated near-term expectations to the incentive structures of the technology sector: “Most of the barriers to this are not technical. The problem is the people driving this are technical. They’re really not business people... They are technicians and they’re massively overestimating because they have to, because they have to get the investors to give them money.” A consulting director drew an explicit historical parallel: “I am old enough to remember when we brought PCs into buildings at a lower level and all of a sudden... we had twenty-five secretaries. And within six months, everyone had been given their own computers, which meant we didn’t need secretaries.” He then noted that the current cycle had not yet reached the analogous inflection point: “What’s going to drive it? What’s the killer application solution that is going to make this? We haven’t got it yet on artificial intelligence.” A non-profit director offered a parallel drawn from the early internet: “I remember going to a workshop many moons ago. Everyone knew how they were using the internet to reach retail customers. But when it came to B2B, people were playing with it at the edges without actually deploying it at scale to deliver business benefit. And in many ways, I draw the same analogy—without having seen what someone else has done, it’s very difficult to reimagine how to do it.”

The Junior Talent Pipeline: A Persistent Structural Concern

A theme that surfaced with consistency across sectors and that had also appeared in earlier stages was the concern that GenAI’s absorption of junior-level tasks would erode the developmental pathways through which expertise is traditionally acquired. Directors framed this not as a speculative future risk but as a dynamic already in motion, and several drew direct

connections to the broader question of why the technology's strategic potential remained unrealized. A policy director articulated the concern as a systemic vulnerability: "AI is very good at getting 80% right. There's 20% grey area, and that you have to refer to someone who really understands it. And the big problem there, if you aren't bringing up sufficient people to come through in 20 years' time and be the experts who can say, well, actually, I remember when there was this serious cancer in that patient here, you're not going to sort out that last 20%." A government advisory director framed the issue in terms of an emerging gap in the organizational hierarchy: "Lower-level jobs—levels one, two, and three, possibly four—are now where AI is. And the human-in-the-loop to five—experience and judgment—is coming in at four, five, six, and above. So, how do you move someone who's relied on AI all this time... if you've never been taught through the experience to actually be able to make those decisions?"

Directors in professional services were particularly attuned to this dynamic. A management consultant observed that the traditional consulting model of training through apprenticeship was being undermined: "What the consultancy is saying is, we don't need so many juniors, because training, onboarding, it's a lot of effort and expense. Some of these junior roles in professional services are easy to replicate by the newest tool." A financial services director who sits on multiple boards described the concern in generational terms: "My worry is people in your generation and below and younger are not going to have the same experience of success and failure to be able to judge. And because AI is backward-looking, my opinion is that the quality of decision-making over time—in two years and five years and 10 years—the quality of decision-making is going to get worse." He continued: "If you don't have a junior analyst on your team that you hired from XYZ University, that is coming up through the ranks... all that work was done by somebody, and that was good work for somebody to build their competencies." An education director pushed the argument further, questioning whether the erosion of junior roles would eventually undermine the very capacity to verify GenAI's output: "As we progress over time, people like me are going to die off and other people who've got that knowledge that they built up through their own experience. And you're going to have a whole caucus of people coming through relying on AI for effectively knowledge and content. And at what point does that human in the loop disappear?" A financial services director who had organized an industry event on the topic captured the emerging consensus: "I'm talking about the looming leadership gap—AI and the potential disappearance of talent and judgment."

Emerging Skepticism and Recalibrated Time Horizons

While directors' strategic expectations for GenAI's eventual competitive implications remained broadly stable, a subset of interviews revealed a shift in the perceived time horizon over which those effects might unfold. Several directors expressed a growing sense that the technology's transformative potential, while real, would take considerably longer to materialize than they or the broader discourse had initially suggested. An independent advisor offered a representative formulation: "When people were talking about it two years ago, it was as if you were saying this software technology on its own, of itself, just because of the nature of what it is and how what it's

capable of doing, is going to be transformative. I find that a bizarre idea, because it always requires people to think about what the impact of that is going to be on the existing environment.” She added: “I just don’t—the software doesn’t do that. It needs to be led, needs to be poked, it needs to be prompted. And, you know, so the entrepreneurial flair of individuals is still necessary.” A consulting director was more blunt: “All the AI companies are saying it’s going to be in two years—rubbish. There’s no way, not enough money and energy in the world to get to that eventuality. We’re talking at least five years.”

Some directors articulated a more fundamental skepticism about the technology’s capacity to move beyond task-level productivity. A consulting director who works across personal productivity, team-level, and enterprise-level applications described a tiered assessment: “You have essentially a relatively constrained technology that is very useful on an individual level. Not particularly useful on—well, it has some use cases—on a team level. And it has almost no use cases on an enterprise level.” He placed the current moment firmly at the beginning of an extended transition: “We’re at the very, very, very beginning of what is going to happen to the economy.” A pharma advisory director described a more personal recalibration, acknowledging that GenAI’s growing capability had paradoxically raised the bar at which clients might bypass his services: “Why should you employ a consultant to provide the expertise when you can get it through GPT? I don’t know the answer to that. And, you know, I might retire and fold up the firm before that happens.” Yet even this more pessimistic assessment was accompanied by the caveat that the effects, when they came, could be profound. A board director in finance, who described GenAI as “way bigger than the internet,” offered a characteristic formulation of this dual perspective: “I think it’s fabulous. I think it’s amazing. I’m a huge fan of AI, I really am. But I’m also, I mean, I’m worried.” A consulting director who works in the brand and advertising sector described competitive effects that were already visible in his industry—“Generative AI has basically destroyed the ad industry—destroyed”—while simultaneously cautioning that the broader economic transition would unfold over decades: “It took probably 20 years from when Britain moved from being a manufacturing-based economy to being a service-based one, and loads of redundancies, and strikes, and basically having to re-educate an entire country... that is what will happen here.”

E2. External Roundtables

In addition to the one-to-one interviews reported above, we conducted two structured external roundtable discussions in early 2026 with participants who had not taken part in the earlier stages of this investigation. Both sessions were conducted under Chatham House Rules to support open dialogue, and participants consented to the disclosure of their identities to provide transparency regarding the composition of each cohort. The first roundtable was held in early 2026 in New York, convened in collaboration with a leading business school; the second was held in early 2026 in San Francisco. Both sessions were designed to revisit the central themes explored throughout the study—task-level adoption, competitive expectations, and the perceived gap between GenAI’s technical capabilities and observable competitive outcomes—with a broader and geographically

distinct sample. These roundtables provided an additional point of triangulation against the longitudinal findings documented in this appendix.

Across both sessions, discussions reflected dynamics broadly consistent with those documented in earlier stages of the study. Participants similarly distinguished between GenAI's task-level capabilities and its competitive implications, and conversations frequently returned to questions of differentiation, displacement, and broader industry structure. While these sessions were not designed to introduce new constructs or materially extend the theoretical framework, they provided additional geographic and institutional variation against which to assess the resonance of the patterns observed in the longitudinal sample.

The New York roundtable was conducted with 23 participants drawn from a similarly senior and cross-sectoral cohort of US-based executives. The anonymized composition of this session is reported below.

Firm	Role / Position	Organization Type
Goldman Sachs	Chief Investment Officer	Investment Banking
Goldman Sachs	Head of AI Accelerator, PE Investments	Investment Banking
Google Research	Distinguished Scientist	Technology (Search and AI)
Microsoft	Economist	Technology (Cloud and AI)
Bloomberg	Head of Responsible AI, CTO Office	Financial Data and Tech
The Estée Lauder Companies	SVP, Enterprise Marketing and Head of AI	Consumer Goods
Procter and Gamble	Director of Strategy and Capability, R&D	Consumer Goods (FMCG)
ExxonMobil	Product Commercial Analytics Manager	Energy (Oil and Gas)
Vodafone	Head of Data Portfolio, Solutions and Innovation	Telecommunications
SAP	Regional Director, Business AI and Tech Strategy	Enterprise Software
Expedia	SVP, Product	Online Travel, Tech
Edward Jones	Partner	Financial Services
AlixPartners	Partner and Managing Director	Management Consulting
Seven2 Private Equity	Partner, Value Creation	Private Equity
State of Massachusetts	Head of AI	State Government
The Conference Board	Center Leader, Marketing and Communications	Business Research
Iterate.ai	Co-Founder and CEO	AI Technology
Evidenza	Co-Founder	AI Technology
BigSpring	Founder and CEO	AI Technology (Workforce)
Business School A	Professor	Academia
Business School A	Professor	Academia
Business School A	Managing Director, Innovation Institute	Academia
Business School B	Researcher	Academia

Note: N = 23. This workshop was conducted under Chatham House Rules. Participants consented to the disclosure of their identities for the purpose of providing transparency regarding the composition of the cohort. Business schools anonymized for peer review.

The San Francisco roundtable brought together 24 participants drawn from both traditional industries and the technology sector, combining senior executives from established firms with founders of AI-native companies, investors, and academics. The composition was intended to stimulate discussion across boundary-spanning perspectives: participants included individuals who build AI products alongside those who adopt them, investors who evaluate AI ventures alongside operators who implement the technology within large organizations, and academics who study innovation and competitive dynamics alongside practitioners who experience them. The table below summarizes the anonymized composition of the San Francisco roundtable.

Firm	Role / Position	Organization Type
Google Research	Distinguished Scientist	Technology (Search and AI)
Qualcomm	SVP and GM, AI Business Unit	Semiconductor Technology
ExxonMobil	Chief Engineer, Data Science	Energy (Oil and Gas)
Procter and Gamble	R&D Vice President, Global Family Care	Consumer Goods (FMCG)
Mondelez Digital Services	Senior Vice President	Consumer Goods (F&B)
Tata Communications	VP and Global Head of New Ventures	Telecommunications
BCG	Partner	Management Consulting
BCG X AI Science Institute	VP, AI and Data Science; Director	Management Consulting (AI)
Fiserv	SVP, Global Digital Solutions	Financial Technology
Fujitsu North America	Regional Chief Technology Officer	IT Services
GoodRx	Chief Product and Technology Officer	Healthcare Technology
BootstrapLabs	Co-Head, Venture Capital and AI Groups	Venture Capital / Private Equity
GrowthCurve Capital	Head of Data, Analytics and ML	Private Equity
GrowthPoint Technology Partners	Managing Director	Technology Investment Banking
Pinegrove Venture Partners	Chief Operating Officer	Venture Capital
SLW	Managing Director	Growth Equity
Evolver	Founder and CEO	AI Technology
ObinAI	Co-Founder and CEO	AI Technology
Emovid	CEO and Co-Founder	AI Technology Start-up
GBK Collective / Mindspan Labs	Founder and CEO; Partner	AI Strategy Consultancy
Martin Reeves Advisory	Founder and Principal	Strategy Advisory
Business School A	Professor	Academia
Business School B	Professor	Academia
Business School B	Research Director	Academia

Note: N = 24. Business school names anonymized for peer review. This workshop was conducted under Chatham House Rules. Participants consented to the disclosure of their identities for the purpose of providing transparency regarding the composition of the cohort. Business schools anonymized for peer review.

Appendix F

Additional Analysis – Matching Stage 2 Survey to Wayback Machine

F1. Methodology & Data Collection

A question that arises from the preceding analysis is whether the heterogeneity in directors' competitive expectations corresponds to heterogeneity in their firms' actual competitive positioning, or whether directors facing substantively similar conditions interpret them through different cognitive lenses. Our roundtable and survey data cannot adjudicate this directly. However, they do permit a preliminary test: if directors' assessments are at least partially grounded in organizational reality, then directors who perceive greater insulation from GenAI should lead firms that are, by independent measure, more distinctively positioned.

To address this, we conducted supplementary analysis examining the relationship between directors' survey responses and an independent, text-based measure of firm differentiation derived from corporate website content, drawing on the approach of Guzman and Li (2023). Specifically, we test whether firms whose public communications exhibit greater competitive differentiation—as captured through natural language processing of website text—are associated with directors who report more positive perceptions of GenAI's potential to enhance their firm's distinctiveness and greater perceived insulation from GenAI eroding their marketplace distinctiveness. This analysis serves two purposes. First, it provides a validity check on directors' self-reported perceptions by examining whether these perceptions correspond to observable features of how their firms describe themselves publicly. Second, it offers preliminary evidence on an empirically underexplored question: whether strategically differentiated firms perceive themselves as better positioned in the face of technological disruption.

We note several important caveats at the outset. The measure captures differentiation in how firms publicly describe themselves, which may diverge from actual competitive positioning or operational distinctiveness. The cross-sectional design precludes causal inference. Additionally, the sample comprises predominantly knowledge-intensive professional services firms, limiting generalizability to other sectors. Nevertheless, this analysis offers an empirical examination of the relationship between competitive positioning and managerial perceptions of GenAI disruption.

We adapt the methodology developed by Guzman and Li (2023) for measuring founding strategy through website text analysis. Their approach uses text-based representations to capture firm positioning, enabling calculation of inter-firm similarity based on the semantic content of corporate communications.

Website Data Collection. For each of the 183 firms in our sample, we collected archived website content from the Internet Archive's Wayback Machine, targeting snapshots from May 2024 to ensure temporal alignment with our survey data collection. Where May 2024 archives were unavailable, we used the nearest available snapshot from 2023–24. Following Guzman and Li

(2023), we extracted (i) the homepage and (ii) up to ten subpages whose URLs contained value-proposition keywords: “about,” “about us,” “company,” “products,” “overview,” “services,” “features,” “solutions,” “what we do,” and “who we are.” This targeted extraction focuses on pages where firms articulate their competitive positioning and value propositions, rather than operational or transactional content. Website text was successfully retrieved for 176 of 183 firms (96.2%).

Text Preprocessing. Raw HTML was converted to plain text through the removal of navigation elements, footers, cookie notices, and other boilerplate content. Text was then preprocessed through lowercasing, removal of URLs and email addresses, removal of standalone numbers, and tokenization. We removed English stopwords and common website boilerplate terms (e.g., “cookies,” “privacy,” “policy,” “copyright”). Documents with fewer than 100 tokens after preprocessing were excluded, yielding a final analytical sample of 161 firms.

TF-IDF Representation. Prior work measuring firm differentiation from web text has employed Doc2Vec embeddings to construct document-level representations (Guzman and Li, 2023). We initially followed this approach but identified a significant confound: Doc2Vec-derived differentiation scores exhibited a substantial correlation with text length ($r = 0.45$), such that firms with longer website descriptions appeared mechanically more differentiated. Because website text length likely reflects firm size, marketing investment, or web design choices rather than strategic positioning per se, this confound threatened the validity of any downstream analysis. This is of lesser concern for the study of startups (Guzman and Li, 2023) but a greater concern for our sample.

We therefore adopted a TF-IDF (term frequency–inverse document frequency) approach, which represents each firm's website text as a sparse vector of weighted term frequencies. Critically, TF-IDF vectors are L2-normalized, which largely eliminates the mechanical relationship between document length and vector magnitude. In our implementation, the correlation between the resulting differentiation score and text length dropped to $r = -0.22$, and we further confirmed that controlling for text length in our regression models did not materially alter the results. We used sublinear term frequency scaling $\log(1 + tf)$, restricted the vocabulary to the 1,000 most informative terms ($\min_{df} = 2$, $\max_{df} = 0.95$), and computed differentiation as each firm's mean cosine distance to its $k = 5$ nearest sector peers — a specification selected through a systematic parameter sweep across ten TF-IDF configurations and six neighborhood sizes.

We represent each firm's website text using Term Frequency–Inverse Document Frequency (TF-IDF) vectors. TF-IDF is a well-established text representation method that weights terms by their importance within a document relative to their frequency across the corpus (Manning, Raghavan, & Schütze, 2008). We retain the 1,000 most informative terms (those appearing in at least 2 but no more than 95% of documents), applying sublinear term frequency scaling to reduce the influence of highly repeated terms. This approach produces a sparse, high-dimensional vector for each firm that captures its distinctive vocabulary relative to other firms in the sample. TF-IDF offers several advantages for our context: it is deterministic (producing identical results on repeated runs),

transparent in its feature construction, and inherently length-normalized through L2 normalization of the resulting vectors.

Differentiation Score Calculation. Following Guzman and Li’s approach (2023), we calculate differentiation as the average cosine distance from a firm’s five nearest neighbors in TF-IDF vector space. Formally, for firm i , the differentiation score is:

$$D_i = \frac{1}{k} \sum_{j \in N_k(i)} (1 - \cos(v_i, v_j))$$

where $N_k(i)$ denotes the set of $k = 5$ firms with highest cosine similarity to firm i , and v_i represents firm i ’s document vector. Higher values indicate greater differentiation—that is, the firm’s positioning is more distant even from its most similar peers. This approach measures relative distinctiveness: how different a firm appears from those firms it most closely resembles. Firms with high differentiation scores occupy unique positions in the competitive landscape, while low differentiation scores indicate clustering with similar firms.

Dependent Variables. Our primary dependent variable is drawn from the survey item: “How do you see Generative AI’s potential to help firms differentiate and create unique value propositions in the market?” (QID26). Responses were recorded on a seven-point scale from “Extremely Negative” (1) to “Extremely Positive” (7). Higher scores indicate that directors see GenAI as more beneficial for their firm’s distinctiveness. As a secondary dependent variable, we examine perceived insulation from GenAI-driven competitive challenges, measured as the average of two reversed survey items: the likelihood that GenAI will challenge (i) the firm’s “distinctiveness in the marketplace: maintaining a unique value proposition amidst widespread Generative AI adoption” and (ii) the firm’s “financial viability: sustaining revenue streams in a market increasingly influenced by Generative AI.” Both items were recorded on five-point scales from “Extremely likely” to “Extremely unlikely” and reversed so that higher scores indicate greater perceived insulation. We also examined a range of other survey items in exploratory analyses.

Estimation. We examine bivariate relationships between text-based differentiation and survey responses using Pearson and Spearman correlations. We then estimate ordinary least squares regression models predicting survey responses from differentiation scores, with and without sector fixed effects. Given the ordinal nature of several dependent variables, we additionally estimate ordered logit models as a robustness check.

Limitations. Several limitations warrant acknowledgment. First, this approach is indirect and limited in discriminating power—website text captures how firms present themselves publicly, not necessarily how they compete in practice. Second, our within-sample comparison measures differentiation relative to other firms in the sample rather than against a broader market; firms appearing differentiated within this sample may be similar when compared to all UK firms. Third,

the methodology assumes that textual similarity maps meaningfully onto competitive similarity—an assumption that, while supported by Guzman and Li's (2023) validation exercises, may not hold equally across all contexts. Fourth, scraping from archived websites introduces potential measurement error where archived snapshots are incomplete or temporally misaligned. Finally, the relatively small sample size ($N = 161$) limits statistical power for detecting modest effects, particularly in models with sector fixed effects.

F2. Descriptive Patterns


Of the 183 firms in our survey sample, 176 (96.2%) had retrievable website content from the Wayback Machine, and 161 met the minimum text threshold of 100 tokens for analysis. The table below presents the sector distribution of the analytical sample.

Sector Distribution of Analytical Sample

Sector	N	Percentage
Professional Services	39	24.2
Other – please specify	21	13.0
Information & Communication	20	12.4
Financial Services	20	12.4
Other	12	7.5
Education	11	6.8
Other Services	10	6.2
Manufacturing	8	5.0
Construction	7	4.3
Real estate	4	2.5
Arts & Entertainment	2	1.2
Health and social work	2	1.2
Electricity and/or gas supply	2	1.2
Mining	1	0.6
Agriculture	1	0.6
Transportation	1	0.6
Total	161	100.0

Note: Sector distribution of the analytical sample ($N = 161$). Sectors are as reported by directors in the survey.

The differentiation score exhibits substantial variation across firms. The mean differentiation score is 0.536 ($SD = 0.143$), with values ranging from 0.125 to 0.793. The interquartile range spans 0.445 to 0.642, and the 10th-to-90th percentile range is 0.360 to 0.719. This distribution suggests meaningful heterogeneity in how distinctively firms position themselves relative to their nearest competitors in the sample.




Distribution of TF-IDF differentiation scores across the analytical sample (N = 161). The dashed line indicates the sample mean.

The mean pairwise cosine distance across all firm pairs is 0.758, corresponding to a mean cosine similarity of 0.242. This indicates that, on average, firms share approximately 24% overlap in their textual positioning. The most differentiated firms tend to operate in specialized technical domains or niche markets, while the least differentiated firms often provide professional services with value propositions common across the sector. This pattern provides face validity for the measure.

As a diagnostic check, we examined the correlation between differentiation scores and text length. The Pearson correlation between differentiation and token count is $r = -0.221$ ($p = 0.005$), indicating a moderate negative association: firms with shorter website texts tend to appear somewhat more differentiated. This is expected given that shorter texts provide fewer shared terms with other documents. However, TF-IDF vectors are inherently L2-normalized, which mitigates direct length dependence. We address this further in our robustness checks by controlling for text length in regression models.

F3. Regression Results

GenAI Distinctiveness Potential. We observe a statistically significant positive correlation between text-based differentiation and directors' perceptions of GenAI's potential to help their firm create unique value propositions. The Pearson correlation coefficient is $r = 0.166$ ($p = 0.036$, $N = 160$). Directors of more textually differentiated firms perceive greater potential for GenAI to enhance their firm's distinctiveness.



Relationship between TF-IDF differentiation and perceived GenAI distinctiveness potential (N = 160). Higher differentiation scores indicate more distinctive textual positioning; higher values on the y-axis indicate more positive perceptions of GenAI for distinctiveness.

The table below presents OLS regression results predicting perceived GenAI distinctiveness potential from text-based differentiation scores.

OLS Regression: Differentiation and GenAI Distinctiveness Potential

	Model 1	Model 2	Model 3
Differentiation Score	1.598 (0.036)	2.101 (0.008)	2.175 (0.008)
Token Count	—	—	0.000 (0.638)
Constant	4.439 (0.000)	—	—
Sector FE	No	Yes	Yes
R ²	0.028	0.167	0.169
Adj. R ²	0.021	0.074	0.069
Observations	160	160	160


Note: Exact p-values in parentheses. Model 2 includes fixed effects for 15 sector categories. Model 3 additionally controls for text length (token count).

In the baseline model without controls (Model 1), the differentiation coefficient is 1.598 (SE = 0.756, $p = 0.036$), indicating that a one-unit increase in differentiation is associated with a 1.60-point increase in perceived GenAI distinctiveness potential on the seven-point scale. Given that the differentiation score's standard deviation is 0.143, a one-standard-deviation increase in differentiation corresponds to approximately 0.23 points higher perceived GenAI distinctiveness potential ($0.143 \times 1.598 = 0.228$), or roughly 17% of a standard deviation in the outcome variable (SD = 1.372).

When sector fixed effects are included (Model 2), the differentiation coefficient increases to 2.101 (SE = 0.786, $p = 0.008$), indicating that the relationship operates within sectors rather than reflecting cross-sector differences in either differentiation or GenAI perceptions. The model R^2 increases from 0.028 to 0.167 with sector fixed effects, indicating that sectors differ in baseline GenAI distinctiveness perceptions, but the within-sector relationship between differentiation and

perceived GenAI potential remains robust. Controlling additionally for text length (Model 3) does not materially alter the differentiation coefficient (2.175, $p = 0.008$), and the text length control itself is not significant ($p = 0.638$), confirming that the result is not driven by text length differences.

Perceived Insulation from GenAI Challenges. We also find a positive association between text-based differentiation and perceived insulation from GenAI-driven competitive challenges. Using the composite measure averaging perceived insulation from distinctiveness erosion and financial viability threats, the Pearson correlation is $r = 0.150$ ($p = 0.058$, $N = 160$), and the Spearman rank correlation is $\rho = 0.155$ ($p = 0.050$). While the Pearson correlation is marginally significant at conventional levels, the Spearman correlation reaches significance at $p = 0.05$, and the direction is consistent with the hypothesis that more differentiated firms perceive greater insulation from GenAI-driven disruption.



Relationship between TF-IDF differentiation and perceived insulation from GenAI challenges (composite of distinctiveness and financial viability insulation, $N = 160$).


OLS Regression: Differentiation and Perceived Insulation from GenAI Challenges

	Model 1	Model 2	Model 3
Differentiation Score	0.980 (0.058)	1.041 (0.056)	1.065 (0.056)
Token Count	—	—	0.000 (0.828)
Constant	2.048 (0.000)	—	—
Sector FE	No	Yes	Yes
R ²	0.023	0.143	0.143
Adj. R ²	0.016	0.047	0.041
Observations	160	160	160

Note: Exact p-values in parentheses. Dependent variable is the average of reversed distinctiveness and financial viability challenge items (1–5 scale). Model 2 includes fixed effects for 15 sector categories. Model 3 additionally controls for text length.

In the baseline model (Model 1), the differentiation coefficient is 0.980 (SE = 0.513, $p = 0.058$). With sector fixed effects (Model 2), the coefficient is 1.041 (SE = 0.541, $p = 0.056$). A one-standard-deviation increase in differentiation corresponds to approximately 0.14 points higher perceived insulation ($0.143 \times 0.980 = 0.140$), or roughly 15% of a standard deviation in the outcome (SD = 0.930). Examining the two insulation items separately, the perceived insulation from threats to financial viability (insulation from “sustaining revenue streams in a market increasingly influenced by Generative AI”) shows a somewhat stronger association with differentiation ($r = 0.128$, $p = 0.054$ with sector fixed effects) than perceived insulation from distinctiveness erosion alone ($r = 0.133$, $p = 0.094$).

The figure below displays the distribution of differentiation scores by perceived insulation category. A general upward trend in median differentiation scores across insulation categories is visible, with firms reporting the highest insulation (category 5) exhibiting higher median differentiation than those reporting the lowest insulation (category 1).



Distribution of TF-IDF differentiation scores by perceived insulation from distinctiveness erosion (N = 160). Categories correspond to reversed responses on the five-point scale (1 = low insulation, 5 = high insulation).

Exploratory Analysis of Additional Survey Items. Beyond our primary dependent variables, we tested the association between text-based differentiation and a broad range of additional survey items to examine whether the pattern extends to other dimensions of directors' perceptions. The table below summarizes these exploratory results, reporting only those items that reached marginal significance ($p < 0.10$) on at least one test.

Exploratory Associations: Differentiation and Additional Survey Items

Survey Item	r	p	Coef (FE)	p (FE)
GenAI distinctiveness potential (1–7)	0.166	0.036	2.101	0.008
Insulation: distinctiveness + financial (1–5)	0.150	0.058	1.041	0.056
Insulation: distinctiveness (1–5)	0.133	0.094	0.850	0.162
Insulation: financial viability (1–5)	0.128	0.106	1.232	0.054
Enhance: workflow processes (1–5)	–0.138	0.081	–0.912	0.118
Enhance: employee productivity (1–5)	–0.144	0.069	–0.970	0.127
Knowledge: relationship (1–5)	–0.192	0.015	–1.121	0.021
Difficulty: competitors overcome barriers (1–5)	–0.150	0.059	–1.144	0.076
Legal liability importance (1–5)	–0.155	0.051	–0.614	0.440
Barrier: compliance & security (1–5)	–0.147	0.063	–1.117	0.104

Note: Table reports survey items where at least one test (Pearson, Spearman, or OLS with sector FE) yielded $p < 0.10$. All models estimated with $N = 160$ –161. Positive r indicates that more differentiated firms score higher on the variable.

Among the noteworthy associations, more differentiated firms report that GenAI will be less likely to enhance workflow processes ($r = -0.138$, $p = 0.081$) and less likely to enhance employee productivity ($r = -0.144$, $p = 0.069$). These findings are consistent with the interpretation that directors of more differentiated firms view their organizations' operations as less susceptible to standardized efficiency improvements through GenAI, potentially because their work is more bespoke or knowledge-intensive.

Directors of more differentiated firms also rate the importance of relationship-based knowledge as lower ($r = -0.192$, $p = 0.015$), which may reflect that these firms' competitive advantage rests more on technical or specialized capabilities than on relationship capital. Additionally, more differentiated firms perceive GenAI compliance and security barriers as less relevant ($r = -0.147$, $p = 0.063$), possibly because their niche positioning involves fewer regulatory constraints or because they have already invested in compliance infrastructure.

We note one result that runs counter to our expectations: more differentiated firms report that it will be easier, not harder, for competitors to overcome their entry barriers using GenAI (difficulty of competitors overcoming barriers: $r = -0.150$, $p = 0.059$). This suggests that while directors of differentiated firms are more optimistic about GenAI's potential to enhance their distinctiveness, they also recognize that the technology may lower traditional competitive barriers.

Robustness Checks. We assessed the robustness of our findings across several dimensions. First, we varied the number of nearest neighbors (k) used to compute differentiation scores. The table below shows that the GenAI distinctiveness result is highly robust across k values from 3 to 20, with the sector-adjusted p-value remaining below 0.05 in all specifications. The insulation composite result also shows consistent directionality, with the Pearson p-value falling below 0.05 for $k \geq 7$.

Robustness: Key Results Across Different k Values

Variable	k	r	p	Coef (FE)	p (FE)
GenAI distinctiveness	3	0.182	0.021	2.622	0.003
	5	0.166	0.036	2.101	0.008
	7	0.157	0.047	2.131	0.012
	10	0.147	0.064	2.137	0.017
	15	0.137	0.084	2.131	0.023
	20	0.128	0.107	2.112	0.030
Insulation composite	3	0.121	0.126	0.801	0.081
	5	0.150	0.058	1.041	0.056
	7	0.157	0.048	1.107	0.046
	10	0.164	0.038	1.163	0.039
	15	0.168	0.034	1.181	0.040
	20	0.170	0.032	1.187	0.041

Note: Table reports Pearson r, Pearson p-value, OLS coefficient with sector FE, and FE p-value for each variable at each k.

Second, we estimated ordered logit models given the ordinal nature of the insulation dependent variable. The ordered logit coefficient for differentiation predicting perceived insulation from distinctiveness erosion is 1.615 ($p = 0.114$), directionally consistent with the OLS results though with wider confidence intervals reflecting the smaller effective sample in each category.

Third, we included text length (token count) as an additional control in our regression models. As reported above, controlling for text length does not alter the GenAI distinctiveness result. For the insulation composite, the differentiation coefficient with both sector fixed effects and text length control is 1.065 ($p = 0.056$), virtually unchanged from the model without the text length control, confirming the result is not an artifact of text length differences.

Interpretation. These results provide some validation that directors' survey responses about their firms' competitive positioning correspond to how the firm describes itself publicly. Firms that communicate more distinctive value propositions on their websites—as captured through TF-IDF analysis of their textual content—are led by directors who see greater potential for GenAI to enhance their firm's distinctiveness and, to a lesser extent, report greater perceived insulation from GenAI-driven competitive threats.

The strongest finding concerns GenAI distinctiveness potential ($r = 0.166$, $p = 0.008$ with sector fixed effects). The coefficient strengthens with sector fixed effects, indicating that the relationship operates within sectors rather than reflecting cross-sector differences. This result is consistent across all robustness specifications and suggests that directors of textually differentiated firms are not simply more complacent about technological change but rather see GenAI as a tool that could complement and enhance their existing competitive positioning. The insulation result, while more modest in magnitude ($r = 0.150$, $p = 0.056$ with sector fixed effects), points in the same direction and is consistent with the interpretation that differentiation may serve as a perceived buffer against GenAI-driven disruption.

We acknowledge that this approach is indirect and limited in discriminating power, and the modest effect sizes—explaining approximately 2–3% of variance in the baseline models—constrain the strength of any conclusions. Nonetheless, the results suggest that at least a portion of the variation in directors' expectations reflects genuine positional differences: some firms really are more distinctively positioned, and their directors' assessments appear to register this. At the same time, the substantial residual variation points toward interpretive processes not straightforwardly anchored in observable firm characteristics—consistent with the cognitive dynamics of threat- and opportunity-framing documented in prior work (Dutton and Jackson, 1987; Jackson and Dutton, 1988). Our data do not allow us to adjudicate the relative contribution of objective positioning versus subjective framing; we can document that both appear to operate, while acknowledging that disentangling them fully exceeds what our design can accomplish.

Appendix G

Additional Analysis – Matching Stage 2 Survey to Fame Firm Data

G1. Methodology & Data Collection

A natural concern with survey-based measures of managerial perceptions is whether they capture substantive variation in firms' competitive positions or merely reflect idiosyncratic optimism or pessimism unrelated to actual firm trajectories. In the main text, we document considerable heterogeneity in how directors perceive the competitive implications of generative AI for their firms and industries. This appendix examines whether these perceptual differences align with objective, independently recorded firm outcomes. Specifically, we link our Stage 2 survey responses to financial and operational data from the FAME database (Bureau van Dijk), which provides standardized audited accounts information for UK companies. Our survey was conducted in early 2024. We examine whether directors' expectations about GenAI's impact, measured at the time of the survey, predict changes in their firms' employment, revenue, profitability, and financial health over the subsequent one to two years. This exercise serves as a form of external validation: if directors' perceptions are informative rather than arbitrary, we would expect them to correlate with realized firm trajectories, particularly given that firms facing greater anticipated disruption may already have been adjusting their operations. We emphasize that this analysis is intended as a validation exercise rather than a causal identification strategy; the survey captures expectations that are partially endogenous to the firm's circumstances. Nevertheless, finding that subjective perceptions predict objective outcomes supports the view that directors' assessments contain real information about their firms' competitive positions in the face of technological change.

FAME data and matching procedure

We obtained financial data from the FAME database for the firms represented in our survey. Our extract covers 168 firms with annual data from 2021 to 2025, encompassing variables including number of employees, turnover, total assets, profit before taxation, shareholders' funds, current ratio, operating profit, EBITDA, profit margin, return on total assets, profit per employee, turnover per employee, fixed assets, intangible assets, and cost of sales. Matching survey respondents to FAME records proceeded in two stages. First, we matched survey responses to our list of 183 IoD member firms using exact company name matching, yielding 187 matched survey responses across 183 unique firms. Second, we matched IoD firms to FAME records using a combination of exact string matching on normalized company names and fuzzy matching (token sort ratio with a threshold of 80), supplemented by 15 manual matches for cases where name variations or trading names prevented automated matching. Four firms were excluded as confirmed mismatches. This process yielded 151 matched IoD–FAME firm pairs, corresponding to 154 survey responses in our final analytical sample.

Dependent variables

For each financial variable, we compute changes over two horizons: 2023 to 2024 (partially concurrent with the survey period, offering larger samples) and 2023 to 2025 (fully post-survey, providing stronger temporal separation but smaller samples). For scale variables such as employees, total assets, shareholders' funds, and turnover, we compute percentage changes as $(\text{Value}(t) - \text{Value}(2023)) / |\text{Value}(2023)| \times 100$, using the absolute value of the baseline to accommodate firms with negative denominators. For ratio variables (current ratio, profit margin, return on total assets), we compute level changes in percentage points. For profit before taxation, which frequently changes sign, we compute absolute changes in thousands of GBP. All dependent variables are winsorized at the 1st and 99th percentiles to mitigate the influence of extreme outliers. Data availability varies across financial variables and years. For the 2023 to 2024 comparison (our primary specification), we have paired observations for 103 firms on employee growth, 123 on total asset growth, 110 on current ratio changes, and 39 on turnover growth. For the 2023 to 2025 comparison, sample sizes are approximately half: 47 firms for employee growth and 53 for current ratio changes, though only 17 for turnover and profitability measures.

Independent variables and composites

Our primary independent variables are the survey perception measures described in the main text. We construct five composite indices as the mean of their constituent items: perceived expected benefit (six items on GenAI's potential to enhance products, build new offerings, improve workflows, boost productivity, enhance creativity, and enable new business models; coded 1–5, higher indicates greater expected benefit); perceived threat (five items on likelihood of challenges to supplier relationships, recruitment, talent training, market distinctiveness, and financial viability; recoded to 1–5, higher indicates greater perceived threat); perceived industry influence (seven items on expected changes to competitive dynamics; coded on a bipolar scale where 1 indicates strongly increase and 5 indicates strongly decrease); competitive imitability (four items on ease of competitor imitation and customer bypass; coded 1–5, higher indicates easier replication); and knowledge importance (five items; coded 1–5, higher indicates greater importance). We also use the displacement score assigned to each firm (1–5, based on sector-level exposure to GenAI disruption), a binary indicator for whether the firm currently uses GenAI, and the respondent's rating of GenAI's importance to their industry (1–5). In multivariate specifications, we control for baseline firm size (log of 2023 employees or log of 2023 total assets) and firm age (log of years since incorporation).

Empirical strategy

We estimate ordinary least squares regressions of the change in an outcome variable on directors' GenAI perception measures, with heteroscedasticity-robust standard errors (HC1) throughout, given the substantial variation in firm size within our sample. We report both bivariate specifications (without controls) and multivariate specifications including baseline firm size and age. We present results for both the one-year (2023–2024) and two-year (2023–2025) horizons, noting that the latter provides cleaner temporal separation from the survey date but at the cost of reduced sample

sizes. The tables below report OLS coefficients with exact p-values in parentheses, following each coefficient. Each cell in the bivariate tables reports a separate regression of the column dependent variable on the row independent variable.

G2. Descriptive Patterns

The table below presents descriptive statistics for the key variables. The mean displacement score is 3.52 (s.d. = 1.08), indicating that the average firm in our sample faces moderate-to-high expected displacement from GenAI. Directors rate GenAI as moderately to highly important for their industry (M = 3.81, s.d. = 1.03), with substantial heterogeneity. The composite indices for expected benefit (M = 3.30) and perceived threat (M = 3.31) suggest that directors on average anticipate both meaningful benefits and meaningful challenges from GenAI. Two-thirds of firms in the sample (66%) report currently using GenAI at the time of the survey. Turning to the FAME financial variables, median employee growth from 2023 to 2024 is zero, with a mean of 8.6% driven by a right tail of rapidly growing firms (s.d. = 40.9). Median turnover growth is slightly negative (−0.4%), and the median change in current ratio is approximately zero. The substantial dispersion in these outcome measures provides meaningful variation for our regression analysis. The sample spans from sole traders to large listed companies: median employment in 2023 is 9 (mean = 827), and median total assets are £699 thousand (mean = £5.0 billion), reflecting the heterogeneous composition of the IoD membership. We control for this size variation in our multivariate specifications.

	N	Mean	S.D.	Min	Max
Panel A: Survey perceptions					
Displacement score (1–5)	154	3.52	1.08	1.00	5.00
GenAI importance (1–5)	154	3.81	1.03	1.00	5.00
Expected benefit index (1–5)	154	3.30	1.00	1.00	5.00
Perceived threat index (1–5)	153	3.31	0.78	1.00	5.00
Competitive imitability (1–5)	153	2.87	0.97	1.00	5.00
Industry influence index (1–5)	153	2.41	0.59	1.00	5.00
Currently uses GenAI (0/1)	149	0.66	0.48	0.00	1.00
Knowledge importance (1–5)	153	4.26	0.69	1.00	5.00
Distinctiveness potential (1–7)	153	5.27	1.42	1.00	7.00
Panel B: FAME financial changes					
Employee growth, % (2023–2024)	103	8.60	40.85	-38.77	360.00
Total asset growth, % (2023–2024)	123	8.22	53.70	-71.89	280.10
Shareholders' funds growth, % (2023–2024)	123	16600.23	183975.81	-1335.79	2040400.00
Current ratio change (2023–2024)	110	0.06	2.96	-14.60	24.20
Turnover growth, % (2023–2024)	39	-0.98	21.56	-55.35	78.11
PBT change, th GBP (2023–2024)	43	-50658.93	362487.35	-2368000.00	122977.60
Return on assets change, pp (2023–2024)	43	-2.56	17.97	-58.31	68.80
Employee growth, % (2023–2025)	47	13.87	31.43	-33.33	100.00
Current ratio change (2023–2025)	53	0.72	6.20	-14.11	24.20
Turnover growth, % (2023–2025)	17	18.99	38.81	-24.32	115.16

Notes: Panel A reports survey perception variables from the Stage 2 IoD director survey (early 2024). Panel B reports changes in FAME financial variables. Percentage changes computed as $(\text{Value}(t) - \text{Value}(2023)) / |\text{Value}(2023)| \times 100$. Level changes reported for ratio variables. All change variables winsorised at the 1st and 99th percentiles. pp = percentage points; th GBP = thousands of pounds sterling.

The table below presents pairwise Pearson correlations between the key perception variables and the primary FAME outcome measures over both time horizons. The pattern of correlations provides initial evidence that directors' perceptions are systematically related to subsequent firm outcomes. Displacement scores and perceived threat are most consistently negatively correlated with employee growth ($r = -0.37$ and $r = -0.41$, respectively, for the one-year horizon). GenAI importance is negatively correlated with employee growth at both horizons and with turnover growth at the two-year horizon. By contrast, GenAI adoption (uses GenAI) shows a positive correlation with current ratio improvement over two years, suggesting that proactive engagement with the technology may be associated with improved liquidity. The correlations between perception variables and turnover or profitability outcomes are generally in the expected direction but often fail to reach conventional significance levels at the one-year horizon, partly because of the smaller sample sizes available for these variables; the two-year correlations tend to be larger in magnitude.

	2023–2024					2023–2025				
	Emp.	Asset	CR	Turn.	ROTA	Emp.	Asset	CR	Turn.	ROTA
Displacement	–0.29 (0.003)	0.15 (0.105)	0.18 (0.057)	–0.12 (0.482)	0.01 (0.967)	–0.20 (0.180)	–0.03 (0.843)	0.32 (0.021)	–0.43 (0.088)	–0.30 (0.243)
Threat	–0.28 (0.005)	0.06 (0.487)	0.14 (0.142)	–0.15 (0.372)	–0.05 (0.771)	–0.31 (0.036)	–0.15 (0.272)	0.22 (0.121)	–0.42 (0.097)	–0.28 (0.277)
Importance	–0.17 (0.079)	0.04 (0.631)	0.19 (0.046)	–0.01 (0.935)	–0.19 (0.226)	–0.32 (0.027)	–0.20 (0.124)	0.32 (0.021)	–0.53 (0.030)	–0.28 (0.282)
Benefit	–0.13 (0.182)	0.08 (0.372)	0.06 (0.526)	0.00 (0.984)	–0.27 (0.080)	–0.17 (0.251)	–0.09 (0.484)	0.18 (0.194)	–0.37 (0.141)	–0.38 (0.138)
Imitability	–0.07 (0.453)	–0.05 (0.559)	0.07 (0.491)	–0.16 (0.329)	–0.14 (0.369)	–0.08 (0.571)	–0.07 (0.583)	0.16 (0.269)	0.16 (0.535)	–0.22 (0.395)
Uses GenAI	0.03 (0.754)	0.06 (0.521)	0.16 (0.108)	0.08 (0.633)	–0.25 (0.113)	–0.17 (0.251)	–0.15 (0.259)	0.33 (0.018)	–0.12 (0.647)	–0.15 (0.580)
Ind. influence	0.12 (0.241)	–0.01 (0.918)	–0.06 (0.560)	0.18 (0.285)	0.02 (0.900)	0.02 (0.905)	0.12 (0.365)	–0.00 (0.994)	0.10 (0.701)	0.14 (0.579)

Notes: Pairwise Pearson correlations with exact p-values in parentheses. Emp. = employee growth; Asset = total asset growth; CR = current ratio change; Turn. = turnover growth; ROTA = return on total assets change. — indicates insufficient paired observations ($N < 10$).

G3. OLS Regressions


Employee growth (2023–2024)

The table below reports regressions of one-year employee growth on directors' GenAI perceptions. The results reveal a consistent pattern: directors who anticipated greater disruption from generative AI presided over firms that subsequently experienced lower employment growth. In bivariate specifications, the displacement score enters with a coefficient of -7.79 ($p = .009$), indicating that a one-unit increase in perceived displacement exposure is associated with 7.8 percentage points lower employee growth. The perceived threat composite is the single strongest bivariate predictor, with a coefficient of -12.22 ($p = .002$), implying that a one-standard-deviation increase in perceived threat corresponds to approximately 9.5 percentage points lower growth. Perceived GenAI importance (-5.18 , $p = .038$), expected benefit (-5.19 , $p = .034$), and competitive

imitability (-3.10 , $p = .045$) are each also significantly negatively associated with employee growth. The negative coefficient on expected benefit is noteworthy: it suggests that directors in firms facing greater competitive pressure from GenAI tend to simultaneously perceive greater potential benefits from the technology—likely reflecting an awareness that adaptation is necessary—while also experiencing headcount reductions as their firms adjust to the new competitive landscape. The individual threat items provide further granularity: directors who reported that GenAI was likely to challenge their firm's distinctiveness in the marketplace saw 7.4 percentage points lower employee growth per unit increase ($p = .011$), while those anticipating challenges to financial viability saw 6.3 percentage points lower growth ($p = .016$).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Bivariate	Bivariate	Bivariate	Bivariate	Bivariate	Bivariate	Bivariate	Full model
Displacement score	-7.793 (0.009)							-3.762 (0.186)
Perceived threat		-12.218 (0.002)						-7.251 (0.076)
GenAI importance			-5.183 (0.038)					-0.127 (0.971)
Expected benefit				-5.192 (0.034)				-2.726 (0.439)
Competitive imitability					-3.104 (0.045)			
Distinctiveness challenged						-7.403 (0.011)		
Financial viability challenged							-6.296 (0.016)	
Uses GenAI								2.674 (0.553)
Log employees (2023)								1.690 (0.159)
Log firm age								-11.479 (0.009)
N	103	103	103	103	103	103	103	101
R ²	0.134	0.169	0.049	0.047	0.020	0.125	0.086	0.279
Adj. R ²	0.126	0.161	0.039	0.037	0.010	0.117	0.077	0.225

Notes: OLS coefficients with exact p-values in parentheses; robust (HC1) standard errors. Dependent variable: employee growth (%), 2023–2024, winsorized at 1st/99th percentiles. Columns (1)–(7) report separate bivariate regressions. Column (8) reports the full multivariate model with controls for baseline firm size and age.



Notes: Bivariate OLS coefficients for predictors of employee growth (%), 2023–2024. Dark bars indicate $p < .05$; light bars indicate $p \geq .05$. Whiskers show 95% confidence intervals.

The final column of the table reports a multivariate specification including the displacement score, expected benefit, GenAI importance, perceived threat, GenAI adoption, and controls for baseline firm size and age. With all perceptions entered jointly, perceived threat retains marginal significance (-7.25 , $p = .076$), and firm age emerges as a strong predictor (-11.48 , $p = .009$), suggesting that younger firms experienced higher employment growth in this period. The joint model explains 27.9% of the variation in employee growth, a substantial share for cross-sectional firm-level data. The figure below provides a visual summary of the bivariate coefficient estimates, with dark bars indicating coefficients significant at the 5% level and lighter bars indicating non-significant estimates; whiskers show 95% confidence intervals. The consistent leftward direction of the significant coefficients underscores the robustness of the finding that greater perceived GenAI disruption is associated with lower subsequent employment growth.

Revenue, profitability, and financial outcomes (2023–2024)

The table below extends the analysis to other financial outcomes over the one-year horizon using multivariate specifications. While sample sizes for turnover and profitability variables are smaller ($N = 37\text{--}43$), several notable patterns emerge. For asset growth ($N = 119$), the displacement score enters positively and significantly (11.63 , $p = .046$), suggesting that firms facing greater expected displacement may have been investing more heavily, potentially in anticipation of needed adaptation. This contrasts with the negative relationship found for employee growth and is consistent with a pattern of capital deepening: firms substituting investment for labor in response to perceived GenAI-driven disruption.

	(1) Asset growth %	(2) Turnover growth %	(3) Current ratio	(4) ROTA change pp
Displacement score	11.632 (0.046)	−3.392 (0.310)	0.285 (0.088)	5.410 (0.287)
Perceived threat	−4.833 (0.623)		−0.184 (0.467)	−4.844 (0.499)
GenAI importance	−7.079 (0.539)		0.350 (0.101)	1.095 (0.659)
Expected benefit	3.897 (0.679)	0.807 (0.819)	−0.366 (0.073)	−4.846 (0.059)
Competitive imitability Uses GenAI	0.556 (0.953)		0.403 (0.203)	−6.916 (0.212)
Log assets (2023)	−0.873 (0.431)		−0.008 (0.853)	
Log turnover (2023)		−0.668 (0.495)		
Log firm age	−9.133 (0.138)		0.350 (0.113)	
Profit pool expectations		5.175 (0.233)		
New revenue expectations		−3.849 (0.307)		
N	119	39	107	42
R ²	0.062	0.065	0.109	0.129
Adj. R ²	0.003	−0.076	0.046	0.008

Notes: OLS coefficients with exact p-values in parentheses; robust (HC1) standard errors. Dependent variables as indicated in column headers, 2023–2024 changes, winsorised at 1st/99th percentiles. Each column reports a single multivariate regression. Blank cells indicate that the variable was not included in that specification.


For the current ratio, displacement score enters with a marginally positive coefficient (0.29, $p = .088$) while expected benefit enters negatively (-0.37 , $p = .073$) in the multivariate specification. The expected benefit coefficient on return on total assets is also negative (-4.59 , $p = .062$ in the bivariate specification), reinforcing the earlier finding that firms whose directors perceived greater GenAI benefits tended to experience worse contemporaneous financial outcomes—possibly because the anticipated benefits had not yet materialized, or because the firms most in need of GenAI-driven transformation are precisely those under greatest competitive pressure.

Two-year post-survey changes (2023–2025)

The table and figure below report results for changes computed over the 2023 to 2025 period. Because these outcomes are measured entirely after the survey date, they provide a cleaner test of whether directors' perceptions predicted subsequent firm trajectories. The trade-off is that sample sizes are substantially smaller, particularly for turnover and profitability measures ($N = 16–17$), and results for these variables should be interpreted with appropriate caution. The most striking finding is the strengthening of the displacement–revenue relationship at the two-year horizon. The displacement score predicts turnover growth with a coefficient of -14.95 ($p < .001$, $N = 17$), and the perceived threat composite enters at -17.97 ($p = .001$). GenAI importance predicts lower turnover growth (-18.65 , $p = .023$) and lower turnover per employee (-10.34 , $p < .001$, $N = 16$). These coefficients are larger in magnitude than the one-year estimates, consistent with the effects of GenAI disruption becoming more pronounced over time. For employee growth over two years ($N = 47$), GenAI importance (-9.89 , $p = .031$) and perceived threat (-11.41 , $p = .067$) predict lower growth, consistent with the one-year results.

	(1)	(2)	(3)	(4)	(5)	(6)
	Employee growth %	Turnover growth %	Turn./emp.	Current ratio	ROTA pp	Profit margin pp
Displacement score	−5.423 (0.256)	−14.947 ($<.001$)	−7.660 (0.008)	1.674 (0.013)	−3.005 (0.056)	−2.009 (0.162)
Perceived threat	−11.410 (0.067)	−17.972 ($<.001$)	−9.609 (0.001)	1.393 (0.129)	−3.399 (0.163)	−1.554 (0.483)
GenAI importance	−9.888 (0.031)	−18.651 (0.022)	−10.340 ($<.001$)	1.878 (0.011)	−2.903 (0.201)	−1.772 (0.351)
Expected benefit	−5.288 (0.221)	−15.107 (0.098)	−7.522 (0.050)	1.063 (0.046)	−4.477 (0.041)	−3.927 (0.025)
Competitive imitability	−2.550 (0.640)	4.826 (0.665)	−8.199 (0.009)	0.817 (0.250)	−2.092 (0.278)	−0.448 (0.783)
Uses GenAI	−10.961 (0.265)	−9.492 (0.654)	−4.735 (0.464)	4.026 (0.007)	−3.116 (0.507)	−5.155 (0.314)
Industry influence	0.834 (0.886)	5.645 (0.652)	−2.359 (0.776)	0.007 (0.991)	2.394 (0.449)	0.561 (0.844)
N	47	17	16	53	17	16

Notes: Each cell reports the coefficient from a separate bivariate OLS regression of the column dependent variable on the row independent variable, with exact p -values in parentheses; robust (HC1) standard errors. All dependent variables measure changes from 2023 to 2025, winsorised at 1st/99th percentiles. N shown is for the displacement score specification; counts may vary slightly across perception variables due to differential item missingness.



Notes: Bivariate OLS coefficients for three key predictors across two-year (2023–2025) outcome variables. Each group of bars represents coefficients from separate regressions. Asterisks indicate $p < .05$. Whiskers show 95% confidence intervals.

The current ratio results are particularly interesting: firms that were already using GenAI at the time of the survey saw a 4.03 point improvement in their current ratio over two years ($p = .007$, $N = 52$), while higher displacement scores (1.67, $p = .014$) and GenAI importance (1.88, $p = .011$) also predict improved current ratios. These patterns may reflect proactive financial management in response to anticipated disruption, or they may indicate that firms with more engaged directors—who both rate GenAI as important and take steps to ensure financial resilience—are systematically different from those whose directors are less engaged. Expected benefit predicts lower profitability outcomes: profit margin change (-3.93 , $p = .025$) and return on total assets change (-4.48 , $p = .042$), reinforcing the interpretation that firms perceiving greater GenAI benefits may be in the midst of costly transitions that have not yet yielded returns. The figure below provides a visual comparison of three key predictors—displacement score, perceived threat, and GenAI importance—across the six two-year outcome variables, with asterisks marking coefficients significant at the 5% level.

Summary and interpretation

The results of this validation exercise support the interpretation that directors' GenAI perceptions, as captured in our Stage 2 survey, contain substantive information about their firms' competitive positions and trajectories. Across multiple specifications and outcome variables, we find that directors who perceived greater disruption from generative AI presided over firms that subsequently experienced lower employment growth, lower revenue growth, and in some cases declining profitability metrics. These relationships are not artefacts of firm size or age: they hold in multivariate specifications that control for baseline firm characteristics. Several patterns merit

particular emphasis. First, the displacement score—a sector-level measure of GenAI exposure—consistently predicts lower employee growth, supporting the view that the occupational displacement framework captures real competitive dynamics rather than abstract technological possibilities. Second, the perceived threat composite is the single strongest predictor of employment changes, with a standardized effect size corresponding to approximately 9.5 percentage points of employee growth per standard deviation. Third, the finding that expected GenAI benefits negatively predict financial outcomes is consistent with a transitional interpretation: firms that perceive the greatest need for GenAI-driven transformation are simultaneously the ones experiencing the most disruptive adjustments.

Fourth, the two-year (2023–2025) results, while based on smaller samples, show generally stronger effects than the one-year results, with the displacement–turnover relationship achieving $p < .001$. This temporal pattern is consistent with the hypothesis that GenAI’s competitive effects are progressive, intensifying as firms and industries adjust to the technology’s implications. The finding that GenAI adopters experienced improved current ratios over the two-year horizon ($p = .007$) suggests that proactive engagement with the technology may confer financial resilience, though the direction of causality cannot be established with our design. We acknowledge several important limitations. Our sample comprises primarily small and medium-sized UK firms, and the results may not generalize to other contexts. The sample sizes for turnover and profitability measures are modest ($N = 16–43$), and the associated results should be regarded as suggestive rather than definitive. Our design cannot definitively establish causality: directors’ perceptions may reflect information about already-occurring changes rather than predictions about future ones. Nevertheless, the consistency of the patterns across multiple outcome measures, perception variables, and time horizons provides meaningful evidence that our survey instruments capture real variation in firms’ exposure to GenAI-driven competitive dynamics.

References

- Brynjolfsson E, Mitchell T. 2017. What can machine learning do? Workforce implications. *Science*. American Association for the Advancement of Science **358**(6370): 1530–1534.
- Charmaz K. 2006. *Constructing grounded theory*, 2nd edition. Introducing qualitative methods. Sage: London; Thousand Oaks, Calif.
- Eloundou T, Manning S, Mishkin P, Rock D. 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science* **384**(6702): 1306–1308.
- Felten E, Raj M, Seamans R. 2021. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal* **42**(12): 2195–2217.
- Flick U. 1998. *An introduction to qualitative research*, 5. ed. Sage: Los Angeles, Calif.
- Gioia DA, Chittipeddi K. 1991. Sensemaking and Sensegiving in Strategic Change Initiation. *Strategic Management Journal*. Wiley **12**(6): 433–448.
- Gioia DA, Corley KG, Hamilton AL. 2013. Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology. *Organizational Research Methods* **16**(1): 15–31.
- Guba EG, Lincoln YS. 1982. Epistemological and Methodological Bases of Naturalistic Inquiry. *Educational Communication and Technology*. Springer **30**(4): 233–252.
- Guzman J, Li A. 2023. Measuring Founding Strategy. *Management Science* **69**(1): 101–118.
- Timmermans S, Tavory I. 2012. *Abductive Analysis: Theorizing Qualitative Research*. University of Chicago Press: Chicago, IL.

Evolution Ltd is a boutique advisory that combines frontier research from world-class business academics and technologists with hands-on experience from senior executives to guide organizations in an increasingly complicated environment.

Evolution focuses on digital ecosystems, Artificial Intelligence and their impact on strategy and organization. Its independence and governance structure ensure rigor and bespoke solutions for its clients and inspire hands-on, [award-winning frameworks](#) that shape managerial practice.

Its clients and partners include large corporates, leading consultancies, governments and NGOs. Projects draw on its affiliates' vast experience and connections to global tech giants, startups, disruptors, entrepreneurs, and governments alike to engage key stakeholders in effective conversations and catalyse action.