



evolution_{ltd}

Research Paper

September 18, 2026

Generative AI and the Architecture of Knowledge Work: Expertise Bypass, Exchange Governability, and Shifting Bottlenecks

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

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

Amartya Das, BCG Henderson Institute, das.amartya@bcg.com

Matt Langione, BCG Henderson Institute, matt.langione@bcg.com

Sesh Iyer, BCG Henderson Institute, sesh.iyer@bcg.com

Contents

3	<u>Executive Summary</u>
4	<u>Abstract</u>
5	<u>Managerial Summary</u>
6	<u>1. Introduction</u>
9	<u>2. Theoretical Background and Emerging Evidence</u>
11	<u>3. Research Design</u>
12	<u>4. Stage 1: Executive Roundtable</u>
12	<u>4.1. Design and Participants</u>
15	<u>4.2. The Unit at Which Participants Located Change</u>
15	<u>4.3. Unevenness Between Capability and Movement</u>
16	<u>4.4. Repositioning and the Vocabulary of Defensible Value</u>
16	<u>4.5. Organizational Types as a Comparative Frame</u>
16	<u>4.6. Summary of Stage 1</u>
18	<u>5. Stage 2: Longitudinal Textual Analysis of Website Language</u>
18	<u>5.1. Data and measurement</u>
19	<u>5.2. Empirical specification</u>
21	<u>5.3. Results</u>
21	<u>5.4. Ex-ante heterogeneity</u>
24	<u>5.5. Summary of Stage 2</u>
25	<u>6. Stage 3: Conditions and Mechanisms of Activity Allocation</u>
25	<u>6.1. Design and Participants</u>
33	<u>6.2. Configuration 1: Contestable Reallocation (High Bypass, Governable Exchange)</u>
34	<u>6.3. Configuration 2: Persistent Bundles, Selective Insourcing (High Bypass, Resistant Exchange)</u>
35	<u>6.4. Configuration 3: Federated Expert Provision (Low Bypass, Governable Exchange)</u>
36	<u>6.5. Configuration 4: Integrated Expert Provision (Low Bypass, Resistant Exchange)</u>
37	<u>6.6. Corollary: Entrepreneurial Positions Opened by the Configurations</u>
39	<u>6.7. Cross-Stage Synthesis: Residual Complements, Shifting Bottlenecks, and Organizational Response</u>
40	<u>7. Discussion</u>
40	<u>7.1. From Task-Level Effects to the Architecture of Knowledge Work</u>
40	<u>7.2. Extending IA: When Production and Exchange Change Separately</u>
41	<u>7.3. Implications for Emerging Organizational Forms</u>
42	<u>7.4. Limitations and Future Research</u>
43	<u>References</u>
48	<u>Appendix</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.

Amartya Das (das.amartya@bcg.com) is a Principal at BCG and an Ambassador at the BCG Henderson Institute, where he leads research on the impact of technology and AI on society.

Matt Langione (matt.langione@bcg.com) is a Managing Director and Partner at BCG, and a leading expert on deep tech commercialization at the BCG Henderson Institute.

Sesh Iyer (sesh.iyer@bcg.com) former North America Chair for BCG X, / Global Leader of the AI & Tech Lab at the BCG Henderson Institute is currently Executive Vice President, Chubb Group and Global Leader, Operations, Technology and Digital Transformation.

Executive Summary

- Most research on Generative AI has asked what the technology can do to tasks. This paper asks a different question: what does it do to the organizations, firms, and industry structures through which knowledge work is produced, bundled, and sold?
- The answer, we find, depends on two dimensions that can move independently of each other. The first is *expertise bypass*: the extent to which GenAI can produce an output without mobilizing the embodied expertise previously required to generate it. The second is *exchange governability*: whether an activity can be coordinated, specified, verified, priced, contracted, and warranted across an organizational boundary. Task-level analyses typically capture the first dimension but not the second. Our research shows why that matters.
- Drawing on a three-stage abductive research design, we combine longitudinal textual analysis of 626 firms observed from 2017 to 2025, a structured roundtable with 27 senior practitioners, and 33 in-depth interviews across two further roundtables, totalling 73 participants drawn from management consulting, financial services, technology, media, healthcare, venture capital and policy.
- Four findings emerge. First, professional services firms have markedly increased their relational language in public positioning since 2022, foregrounding client relationships, advisory judgment, trust, and accountability, while interface-oriented technology firms have increased their emphasis on scalable systems. Second, qualitative accounts reveal uneven shifts in knowledge-work activities: some have moved from bundled engagements to technology platforms, client organisations, or teams of independent specialists, while others have been transformed in place and others remain stable. Third, crossing the two dimensions of expertise bypass and exchange governability yields four configurations that organize these patterns of movement, transformation, and stability. Fourth, two new organizational forms are emerging around the off-diagonal configurations: acquire-and-automate consolidators that absorb incumbent service businesses and automate their administrative core, and federated specialist teams enabled by newly governable interfaces.
- The paper contributes to two bodies of literature. It shows why moving directly from task-level capability to organizational outcome is incomplete: a production gain does not by itself determine whether an activity stays inside an incumbent, moves to a client or platform, or supports new entry. That also depends on exchange governability and the residual complements firms control. It also extends Industry Architecture theory by making expertise bypass an explicit and independently varying production dimension of architectural change, identifying persistent bundles with selective insourcing and federated expert provision as outcomes that task capability alone cannot predict.
- For organizations, the key implication is that GenAI exposure is not a property of a firm's tasks alone. It is a property of how those tasks are embedded in workflows, accountability structures, and the conditions under which outputs can be verified and warranted. Understanding which activities are genuinely contestable, and which remain bundled by liability, accreditation, or relational complexity, is the first step toward a more rigorous strategic response.

Abstract

Research on generative AI (GenAI) and work has largely focused on tasks. We ask when task-level effects alter knowledge-work industry architecture: the allocation of activities across firm boundaries. Using a three-stage abductive design - an executive roundtable, a panel of 626 firm websites (2017–2025), and 33 interviews plus two further roundtables with C-level executives and corporate directors - we document divergent repositioning: professional services firms foreground relational value, while interface-oriented technology firms emphasize scalable interfaces. We find that activity reallocation is shaped by two independently varying dimensions: expertise bypass and exchange governability. Their intersection yields four configurations: contestable reallocation, persistent bundles with selective insourcing, federated expert provision, and integrated expert provision. We extend industry architecture theory by explaining change when production and exchange conditions shift separately.

Managerial Summary

GenAI is not simply making knowledge work more efficient - it may be beginning to change which organizations perform which work. We document that professional services firms are emphasizing relationship-driven, judgment-intensive work, while interface-oriented technology firms are extending scalable delivery. We develop a framework explaining why task exposure alone does not determine where work moves. We distinguish between expertise bypass - whether GenAI can produce an output without the expertise previously required — and exchange governability — whether that activity can be specified, verified, priced, and warranted across organizational boundaries. Where both are present, work becomes contestable and may be reallocated. Where they diverge, highly exposed activities may remain bundled, while less exposed, human-intensive work may move through federated specialist models. These conditions shape which organizational forms become viable.

Keywords: generative AI, industry architecture, repositioning, knowledge work, pre-adaptation

1. Introduction

Generative AI has been assessed mainly through what it does to tasks. Field studies ask whether workers complete writing, coding, analysis, consulting, and related tasks faster or better; exposure studies map which tasks are technically susceptible; labor-market studies track shifts in demand (Brynjolfsson, Li, and Raymond, 2025; Dell'Acqua et al., 2026; Demirci, Hannane, and Zhu, 2025; Eloundou et al., 2024; Massenkoff and McCrory, 2026; Noy and Zhang, 2023; Teutloff et al., 2025). These capabilities bear most directly on knowledge work, whose tasks consist largely of producing, evaluating, and communicating information. It is increasingly clear what GenAI can do to such tasks. It is much less clear how those task-level changes reorganize the division of labor through which knowledge work is produced, bundled, and sold.

Knowledge work reaches the market embedded in organizations, contracts, and claims to expertise (von Nordenflycht, 2010; Susskind and Susskind, 2015). The value of an analysis or a legal opinion depends on how the problem was framed, on the data it draws on, on who will rely on it, and on who is accountable for it (Abbott, 1988). GenAI can make expert output easier to produce without making the associated activity easier to trade (Chatterji, Rock, and Talamàs, 2026; Garicano, Li, and Wu, 2026). A given task-level change can therefore be absorbed in several ways: an incumbent may perform the work with fewer people or at lower prices, the work may move across a firm boundary in either direction, or it may be taken up by entrants.

Therefore, to understand how GenAI affects the division of labor between different kinds of firms, we have to consider both conditions of exchange (Williamson, 1985) and relative productive capabilities of different firms (Jacobides and Winter, 2005). Industry Architecture (IA) research examines how activities move across firm boundaries when coordination becomes manageable and outputs can be specified and verified across an interface; firms then compete to shape those interfaces and control the bottlenecks and complements that remain (Baldwin, 2008; Hannah and Eisenhardt, 2018; Jacobides, Knudsen, and Augier, 2006; Kapoor and Raj, 2026; Teece, 2026). GenAI is unusual because it can change both production and exchange. It may therefore do more than automate tasks or reduce labor: it may redraw organizational boundaries and help us extend IA theory itself.

A substantial body of research documents AI tools performing work long considered dependent on tacit human expertise (Autor, 2015; Brynjolfsson and Mitchell, 2017; Choudhury, Starr, and Agarwal, 2020; Krakowski, Luger, and Raisch, 2023), and large language models extend this reach to the text, code, and analysis that professionals sell (Eloundou et al., 2024; Noy and Zhang, 2023). GenAI therefore acts directly on the production side of IA by changing whether outputs can be generated without mobilizing the firm-specific, embodied capabilities that previously produced them (Jacobides and Winter, 2005, 2012). It may also affect exchange, but comparability, verification, accountability, contracts, accreditation, and regulation often change on different timescales (Marchant, 2011; Williamson, 2000). We ask how changes in production and exchange together shape the allocation of knowledge work.

We investigate this through a three-stage abductive design (Timmermans and Tavory, 2012), observing the setting early, before a new architecture has settled. Stage 1 uses a structured roundtable to establish how practitioners categorize the changes underway. Stage 2 follows the public positioning of 626 firms from 2017 to 2025 and investigates whether organizational types

repositioned differently after ChatGPT; website language indicates where firms claim differentiation. Stage 3 draws on 33 interviews and two further roundtables and moves beneath the firm to the activity, asking how each piece of knowledge work is produced, specified, verified, and paid for. Accounts are coded as realized, reversed, prospective, or non-moving, and claims of reallocation rest on realized accounts with identifiable referents.

Three findings emerge. First, we document a post-ChatGPT trend break in website language: professional services firms increasingly emphasize client relationships, judgment, and trust relative to traditional corporations. Subsector analysis identifies increased emphasis on scalable digital systems among software, platform, and AI/data firms. To understand how these shifts in positioning relate to the organization of work, we examine individual activities. Second, interviews and roundtables identify two dimensions that help explain activity reallocation, or its absence: whether expert outputs can be produced without mobilizing the embodied expertise previously required (expertise bypass), and whether the activity can be coordinated and traded across firms with verifiable outputs and clear accountability (exchange governability). These dimensions can vary separately, yielding four configurations. Third, informants describe emerging or prospective entrants, especially around the off-diagonal configurations, including acquire-and-automate consolidators and federated specialist teams. Taken together, these findings move beyond a task-level account to identify an emerging logic behind the reallocation and transformation of knowledge work within and across firms.

The paper contributes to two discussions. First, research on AI and work typically moves from task capability to organizational outcomes. We show why that move is incomplete. A task-level production gain matters, but it does not by itself determine whether an activity stays inside an incumbent, moves to a client or platform, or supports entry. That also depends on exchange governability and on the residual complements firms control.

Second, we extend IA by specifying how expertise bypass interacts with exchange conditions to shape architectural change. IA already treats productive capabilities and exchange conditions as jointly shaping firm boundaries, but its empirical applications have generally examined settings in which producing the focal output still required the embodied expertise of people and firms. GenAI relaxes that condition: an output may now be generated without mobilizing the expertise previously required, while the exchange around it may remain unchanged. The off-diagonal cases make the contribution visible. Legal analysis and audit can be highly exposed to GenAI yet remain bundled because liability, accreditation, and accountable sign-off resist separation. Construction design expertise can remain human yet be redistributed across federated specialists once shared interfaces make the exchange governable. Federated expert provision recovers IA's established account of exchange-driven separation; persistent bundles with selective insourcing show how expertise bypass can transform production while separation remains constrained. More broadly, the setting connects two concerns at the core of IA: how activities are divided across firms and how firms seek to appropriate value within the resulting architecture. Where GenAI makes expert production cheaper and more widely available, it changes what remains scarce and therefore the bundles and interfaces around which organizations can defend value (Zhang, Wuebker, and Barney, 2026).

Because the setting is early and unsettled, we keep the claims narrow. Website language measures firms' stated value propositions, not boundary change (following Guzman and Li, 2023). We claim reallocation only where informants described realized changes with identifiable

referents, and we do not argue that GenAI has already transformed the organization of knowledge work. Our claim is that the early movements are patterned: they turn on what stays bundled, what becomes separable, who can govern the newly separable activity, and who controls the complements that become more valuable as generation becomes cheap.

2. Theoretical Background and Emerging Evidence

Most empirical work on GenAI and knowledge work remains at the level of the task or worker. Field studies document sizeable effects on productivity and quality in customer support (Brynjolfsson et al., 2025), consulting (Dell'Acqua et al., 2026), idea generation (Boussiou et al., 2024; Girotra et al., 2023), software development (Gambacorta et al., 2024; Hoffmann et al., 2025; Peng et al., 2023), and skill acquisition (Wiles et al., 2024); exposure indices rank knowledge-work activities highly (Eloundou et al., 2024; Handa et al., 2025); and online labor-market studies trace shifts in demand (Demirci et al., 2025; Teutloff et al., 2025; Yiu et al., 2025). These designs are well suited to showing what GenAI does to a task. For that purpose, they usually hold the task's organizational location fixed - whether it sits in the billable core of a professional services firm, an internal corporate function, a specialist supplier, or a platform product. When that location changes is a different question, and it is the one we pursue.

IA examines such changes. An industry's architecture is its sector-wide division of labor: the allocation of activities across firm boundaries and the rules and interfaces that sustain it (Jacobides, Knudsen, and Augier, 2006). The lineage runs from the division of labor and the extent of the market (Smith, 1776; Stigler, 1951), through transaction-cost reasoning about when exchange is feasible (Coase, 1937; Williamson, 1975, 1985), to work on modularity showing that transactions concentrate at "thin crossing points" where limited information must pass between activities and can be standardized and contracted over (Baldwin and Clark, 2000; Baldwin, 2008). Architectures evolve endogenously as firms pursue gains from trade, first partitioning activities within organizations and then across them, while capabilities and transactional feasibility co-evolve (Jacobides, 2005; Jacobides and Winter, 2005). Once an architecture settles, firms compete to shape exchange conditions, occupy bottlenecks, and commoditize the activities of others (Jacobides, Knudsen, and Augier, 2006; Jacobides and Tae, 2015; Hannah and Eisenhardt, 2018).

IA already treats productive heterogeneity and transaction costs as jointly shaping firm boundaries. Its architectural predictions, however, have typically turned on when interfaces become governable: coordination between adjacent activities must become manageable, and what crosses the boundary must be standardized enough to be specified *ex ante* and verified *ex post* (Jacobides, 2005; Jacobides and Winter, 2005; Thompson, 1967; Kogut and Zander, 1992). Together these conditions determine the governability of an activity's exchange - whether it can be coordinated, specified, verified, priced, contracted, and warranted across an organizational boundary. We assess governability relative to the proposed separation of the focal activity from its existing bundle. Resistant exchange therefore denotes difficulty in establishing a new interface for its provision. In the empirical settings most often studied, the capability to produce the focal output remained embodied in people and firms. Mortgage origination, for example, became tradable when standardized loan characteristics and credit scores made the exchange governable, while the originator's expertise remained tacit and human (Jacobides, 2005).

IA as a framework also helps explain the comparative stability of knowledge work. Three features reinforced one another. First, the expertise producing the work was substantially tacit: professionals could not fully articulate what they knew (Polanyi, 1966; Nelson and Winter, 1982), and codifying it was costly, partial, and slow (Zander and Kogut, 1995; Cowan, David,

and Foray, 2000), so production could not be separated from the expert. Second, the outputs were difficult to evaluate: much knowledge work has the character of a credence good, whose quality buyers cannot readily assess even after delivery (Darby and Karni, 1973; Dulleck and Kerschbamer, 2006), a problem professional services research identifies as opacity (Sharma, 1997; von Nordenflycht, 2010). Third, in place of verification, exchange rested on institutional substitutes - professional jurisdiction, accreditation, reputation, and liability - that tie the reliability of an output to an accountable expert or firm (Abbott, 1988). The resulting division of labor was organized around professions, bundled engagements, and knowledge hierarchies (Garicano, 2000; Garicano and Rossi-Hansberg, 2006), largely because expertise remained embodied in people.

Earlier information technologies entered this architecture along a codification route. Computerization substituted for tasks that could be expressed in explicit rules (Autor, Levy, and Murnane, 2003), and enterprise systems automated work after the adopting organization had rendered it explicit in procedures and structured data (Zuboff, 1988); the tacit remainder defined the technological frontier (Autor, 2015). GenAI departs from this route. Trained on large corpora of human-produced output, large language models internalize regularities in expert work, including patterns their producers could not themselves articulate (Brynjolfsson, Li, and Raymond, 2025), and generate analyses, drafts, and code that resemble expert output. The focal organization does not first formalize the underlying expertise into task-specific rules. Expertise bypass allows an actor to produce an expert output without acquiring or mobilizing the embodied expertise previously required for that production step. Framing, verification, and accountability may continue to require expertise elsewhere in the workflow. Because these models operate in natural language, the medium in which professional work is specified, delegated, and reviewed, the capability reaches many domains at once, a characteristic of general-purpose technologies (Bresnahan and Trajtenberg, 1995). Model-generated output does not itself supply an accountable warrant; review and responsibility must be provided through the surrounding organizational arrangements. Reliability is hardest to assess precisely where the credence problem was already most severe (Anthony, 2021; Gennai et al., 2026; Lebovitz, Levina, and Lifshitz-Assaf, 2021).

GenAI therefore creates a distinctive theoretical problem. It can change whether an output is produced without mobilizing the embodied expertise formerly required, while the conditions under which that output can be verified, warranted, and exchanged may change slowly or not at all. Conversely, shared interfaces can make human expertise easier to trade even where producing the output continues to require specialists. The production and exchange dimensions can therefore change separately. GenAI may also lower coordination costs by translating across heterogeneous systems without requiring prior standardization (Choudary, 2025), or make some outputs more demonstrable. Yet exchange still depends on contracts, accreditation, professional jurisdiction, liability, procurement rules, and other institutions that develop on their own timescales (Goldfarb and Tucker, 2019). We investigate how the two dimensions combine to shape activity allocation. Because the setting is still unfolding, changes are uneven across activities, and it is too soon to know where they will ultimately settle, our analysis is explicitly abductive: it seeks to identify emerging mechanisms and configurations rather than predict a settled architecture.

3. Research Design

Because GenAI's effects on knowledge-work industries are still unfolding and their organizational implications remain theoretically unclear, we chose an abductive research design (Pillai et al., 2020; Timmermans and Tavory, 2012) suited to developing explanations for phenomena that existing theory does not yet adequately explain (Choudhury, Moeen, and Wormald, 2026), where deductive hypothesis testing would be premature. The design combines purposeful sampling and complementary evidence across three stages (Bamberger, 2018; Creswell and Plano Clark, 2018), each informing the next (Behfar and Okhuysen, 2018). Each stage answers a different question. Stage 1 provides the practitioner vocabulary and four organizational types used in Stage 2. Stage 2 examines whether those types repositioned differently in public after GenAI's introduction. Stage 3 moves to the activity level, examining realized and prospective changes in allocation and developing the framework that organizes them. Stage 2 builds directly on Stage 1; Stage 3 develops related but distinct constructs. We therefore do not treat Stage 3 as validating the Stage 2 framework. The stages build an explanation rather than independently measuring one construct. Stage 2 supports claims about public positioning. Stage 3 supports claims about perceived changes in conditions and, where accounts have identifiable referents, early realized changes in allocation and delivery. Only realized accounts support claims about boundary movement. We reserve reallocation for those accounts and use architectural pressure for the broader pattern. The Online Appendix reports coding procedures, data structures, and disconfirming evidence.

4. Stage 1: Executive Roundtable

4.1 Design and Participants

In June 2025 we convened a structured roundtable of 27 senior practitioners drawn from management consulting, financial services, technology, media, logistics, healthcare, venture capital, and policy, each with direct engagement in GenAI adoption decisions (Miles and Huberman, 1994); Table 1 details the composition.

Table 1: Stage 1 — Roundtable Composition

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
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
Boutique Advisory	Engagement Manager	Leadership Advisory
Business School	Professor; Lead Advisor	Academia
Boutique Advisory	Director	Management Consulting

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.

The session comprised a plenary, two concurrent breakout groups, and a reconvened plenary; both rooms were recorded and transcribed verbatim. Discussion was deliberately open-ended, organized around how participants expected GenAI to change knowledge work in their sectors and what they had already observed, and we pressed on the difference throughout. The accounts range from expectations, through decisions taken but not yet implemented, to changes reported as already in place, and we mark the distinction where it bears on a claim. Analysis followed established protocols for inductive research (Eisenhardt, 1989; Gioia, Corley, and Hamilton, 2013), with independent coding reconciled through iterative discussion; the Online Appendix reports the data structure and disconfirming accounts. Because the roundtable preceded any framework of our own, its role is diagnostic: it establishes the unit at which participants located change, the unevenness between what the technology could produce and what had been reorganized, and the vocabulary in which organizations described their responses. The first two motivate the activity-level analysis of Stage 3; the third supplies the measurement framework for Stage 2. Table 2 provides representative quotations by theme.

Table 2: Stage 1 — Representative Qualitative Evidence by Empirical Theme

PANEL A: THE UNIT AT WHICH PARTICIPANTS LOCATED CHANGE	
Decomposition fails at the joins between tasks	“A lot of project leader work is in between tasks. It is really assigning tasks, managing tasks, getting the response and connecting things together.” – Partner, consulting “Are your scaling tasks and your specialization tasks co-located within one role? ... Can they be separated out?” – Chief data officer, technology
Cheap generation displaces effort onto verification	“You need to know how to verify the output ... that is a job of the project leader, who looks at the results coming from associates and consultants and says, ‘well, that doesn’t make sense.’” – Partner, management consulting “You can produce a lot of just good enough content really, really quickly, at significant volume, and you can swamp some of your colleagues who produce much better quality.” – Partner, leadership advisory
Uptake follows the workflow, not the capability	“If we build a tool that embeds in one of those platforms, then the adoption rate overnight goes from zero to 100 ... The horizontal stuff is much, much more difficult.” – Global managing editor, news agency “Your partner does not need to know how to make slides with ChatGPT ... What they want to know is that first 10 minute conversation with the CEO.” – Partner, management consulting
PANEL B: UNEVENNESS BETWEEN CAPABILITY AND MOVEMENT	
Capability present, reallocation absent	“[The firm] has essentially automated the process of audit for [sector] ... we can now autonomously conduct audits better than a human being.” – Chair, science and innovation advisory “Ask them how many associates they’re still hiring today ... they don’t need anyone to do research.” – General partner, venture capital
The binding constraint is institutional rather than technical	“Job displacement is inversely correlated with the level of control, regulation and qualifications required for the job.” – Head of AI, insurance “In legal AI, it’s easy to sell to GCs – general counsels – rather than to law firms ... But then in law firms, there’s just resistance.” – Advisor, venture capital “Where I saw there was a difficulty in adopting GenAI is where you have a part of the organization that is around a mastery mindset ... people tended to feel that AI in general was like cheating.” – Director, human resources strategy
The terms on which an activity can be transacted vary with scale	“We have an agreement ... that none of our data will be used to train any future models ... We can do that because we are a huge global organization ... Would a mid-cap firm be able to have that negotiating power?” – Research vice president, industry research “The boards of those SMEs are not able to absorb into their governance framework how to govern this thing.” – Chair, science and innovation advisory
PANEL C: REPOSITIONING AND THE VOCABULARY OF DEFENSIBLE VALUE	
Locating the boundary of commoditization	“What we see commoditizing really, really quickly is a lot of the why and the what. What we don’t see commoditizing anytime soon is more oriented around the how and who.” – Partner, leadership advisory
Repositioning within the existing bundle	“For executing the same job next year, you have to be 10 per cent less expensive, then you don’t have any other choice [but] adapting your company model.” – CFO, management consulting “Demand stays fixed, people work faster – so we need to shrink ... that changes the entire foundation of this up-or-out model.” – Partner, management consulting
Disconfirming account: demand absorbs the gain	“We have more junior people per senior people today than we did 15 years ago, because client expectations on volume and quality and analysis have gone through the roof faster than the technology has been able to keep up.” – Partner, leadership advisory
Technology-side repositioning runs through distribution	“It actually largely doesn’t go into compute. [It] goes to go-to-market ... the friction of actually trying to get AI technology into large companies is huge.” – Advisor, VC “The differentiation between an okay software company and a great software company has absolutely nothing to do with the software. [It] has to do with the sales machine.” – General partner, venture capital

Note: n = 27. Brackets indicate inserted or anonymized text; ellipses indicate omissions. Quotations identify role and sector only. Appendix A reports the full evidence and coding details.

4.2 The Unit at Which Participants Located Change

Participants accepted the task-level account of what GenAI can do while disputing its sufficiency for what they were seeing. A consultancy that had tried to have software assume a middle-management role reported that the attempt failed at the joins between tasks: "a lot of project leader work is in between tasks. It is really assigning tasks, managing tasks... connecting things together" (Partner, management consulting). Adoption inside the same firm concentrated among associates, whose work exposure indices rank as most susceptible, and thinned at the layer above. A chief data officer recast the categories accordingly, asking whether "scaling tasks" and "specialization tasks" must remain co-located within one role (Chief data officer, technology). Where production had changed substantially, participants separated generation from use and described growing demands for verification: responsibility moved upward to whoever was accountable for the output ("you actually need to know how to verify the output... that is a job of the project leader"; Partner, management consulting), while a competent user could "produce a lot of just good enough content... at significant volume" and swamp colleagues producing better work (Partner, leadership advisory). Uptake followed the workflow rather than capability, moving fastest where tools were embedded in platforms staff already used (Global managing editor, news agency). Participants therefore located change at the level of an activity embedded in a workflow and accountability structure: "what we see commoditizing quickly is a lot of the why and the what. What we don't see commoditizing anytime soon is more oriented around the how and who" (Partner, leadership advisory).

4.3 Unevenness Between Capability and Movement

Participants' accounts of what had changed did not track their accounts of what the technology could do. An entrepreneur reported having "essentially automated the process of audit" for cybersecurity compliance, yet none of the participants reported a corresponding reorganization of an audit firm. Associate hiring in law was described as already substantially reduced while the production and signing of advice remained unchanged, and participants attributed the gap to accountability: "job displacement is inversely correlated with the level of control, regulation and qualifications required for the job" (Head of AI, insurance). The same capability produced different outcomes in differently structured buyers: legal AI sold readily to general counsel and met resistance in law firms (Advisor, venture capital), while pharmaceutical R&D budgets, set by therapeutic area, left horizontal technologies without a budget line that technology buyers organized by capability could supply (VP, industry research). The terms on which an activity could be transacted also varied with scale. One large firm had secured an agreement that none of its data would train future models and doubted a mid-cap firm could obtain the same terms; smaller organizations were described as constrained by liability and governance rather than capability (Chair, science and innovation advisory). Where acceptable terms could not be obtained, use moved out of view rather than stopping, with formally banned public tools accessed "from their phones" (Director, boutique advisory). Effort also shifted toward the conditions of reliance: in corporate and public-sector deployments, "the security, the accountability, the auditability, the explainability... is actually a majority of the effort" (Director, applied AI consultancy). These accounts show production and the arrangements through which outputs are bought, verified, and relied upon changing at different rates and for different reasons; Stage 3 returns to that separation activity by activity.

4.4 Repositioning and the Vocabulary of Defensible Value

Participants discussed their organizations' responses at the level of what the organization sells and described shifts in emphasis within an existing bundle. A news agency reported that its most exposed newsroom had become its most enthusiastic adopter, with the savings reinvested "from that speed and scale to the higher end" (Global managing editor, news agency). Price pressure appeared in clients' refusal to pay for newly abundant creative work (CFO, management consulting) and in the leverage structure of professional firms, whose exposed activities sit at the base of the pyramid: "demand stays fixed, people work faster - so we need to shrink... that changes the entire foundation of the business model of the consulting world" (Partner, management consulting). A disconfirming account came from the same group of firms: a leadership advisory partner reported higher junior-to-senior ratios than fifteen years ago because client expectations had absorbed prior productivity gains. Technology participants described repositioning around distribution and specificity, with venture funding going "to go-to-market... to salespeople" (Advisor, venture capital) and model providers moving toward function- and industry-specific applications. Across organizational types, participants converged on a recurring vocabulary: relational work, covering client advisory, trust-building, liability-bearing, and institutional legitimacy; and executional work, covering the analytical, procedural, and administrative activities within GenAI's immediate reach. We use these as broad practitioner-derived labels rather than theoretical constructs. The characteristics grouped under each label need not go together, and participants described cases that cut across the distinction, including resistance to GenAI among staff with "a mastery mindset" (Director, HR strategy).

4.5 Organizational Types as a Comparative Frame

Participants grouped the organizations performing knowledge work into four recurring types: professional services firms, bundling executional and relational work within a leveraged engagement; traditional corporations, carrying legacy processes alongside institutional experience of prior transitions; incumbent technology firms, integrating AI into existing product suites; and emerging technology firms, competing on specialization and speed. The additions they proposed were informative: small and medium enterprises as a distinct type on liability and governance grounds, firm age as a sharper cut than sector, funded organizations as a category of their own, and an AI-native, service-driven roll-up that acquires an incumbent service business, keeps its client relationships, and automates its administrative component. We retain the four types for Stage 2, where a consistent longitudinal sample can be assembled.

4.6 Summary of Stage 1

The roundtable produced three provisional outputs. Participants located change at the level of an activity embedded in a workflow and accountability structure, identifying the joining and checking of tasks as what had not moved. They reported a gap between what the technology could produce and what had been reorganized, attributing it to regulation and professional accountability, the structure and scale of the buyer, and the terms on which an activity could be

transacted. They also described repositioning within an existing bundle, using relational and executional language to identify defensible value. These are practitioner accounts from 27 individuals, each speaking from a particular organization and sector. Whether this repositioning vocabulary appears systematically in firms' public claims is the narrower question to which we now turn. The activity-level allocation question remains for Stage 3.

5. Stage 2: Longitudinal Textual Analysis of Website Language

Stage 2 asks whether the repositioning vocabulary from Stage 1 appears systematically in a larger longitudinal sample. In a loose analogy, the four organizational types are different "species" exposed to the same GenAI shock, but with different productive capabilities, exchange conditions, and inherited complements. Website text provides a public record of how they present their basis of differentiation. It does not reveal work allocation, internal redesign, or boundary change, so we treat it as evidence of public positioning, not reallocation (following Guzman and Li, 2023).

That evidence is still useful in an abductive design because positioning is an early, observable response to architectural uncertainty. If GenAI changes which components of knowledge work appear contestable, firms may adjust the language through which they claim differentiation. We therefore ask whether relational and executional language changed systematically after the launch of ChatGPT in late 2022, whether those changes departed from pre-existing trajectories, and whether they were concentrated in subsectors whose underlying work differed *ex ante*.

5.1 Data and measurement

We construct a longitudinal panel from archived firm websites. For each firm-year, we retrieve the homepage and up to ten keyword-filtered subpages from the Wayback Machine, clean archive boilerplate, and retain text that passes minimum coverage screens. The analytic panel contains 626 firms observed over 2017-2025, yielding 5,330 firm-year observations: professional services firms (168), emerging technology firms (115), incumbent technology firms (126), and traditional corporations (217). Because archival coverage is thinner farther from the present, we also define a balanced pre-period sample of 453 firms observed in each of 2017, 2018, and 2019.

Following Hassan et al. (2025), we construct purpose-built dictionaries to measure the relational and executional vocabularies that emerged from Stage 1. The primary relational measure is a 50-term dictionary capturing claims around client relationships, advisory expertise, partnership, trust, contextual understanding, stakeholder responsibility, and values-based accountability; two constituent measures, a client/advisory dictionary and a stakeholder-responsibility dictionary, assess whether results depend on a single broad construction. Executional language captures claims around automation, deployment, workflow, delivery, implementation, scale, and operational efficiency, likewise measured with a comprehensive dictionary, an interface/scalable-systems dictionary, and a delivery/implementation dictionary. Terms are matched using word boundaries on lowercased text and normalized as counts per 10,000 words. The Online Appendix reports the term lists, measurement-validation checks, and validation against Harvard IV-4 General Inquirer categories (Stone et al., 1966).

The Online Appendix reports specification ladders and GenAI-mention controls for the extended panel, term-level measurement diagnostics, and supplementary Harvard IV-4, placebo-dictionary, permutation, and expanded-sample checks for the 2020-2025 post-period models. These checks address different aspects of measurement and sample sensitivity; the latter

exercises do not directly test the post-2022 slope-break specification. The substantive interpretation remains tied to differences in public positioning rather than to a causal estimate of technology adoption.

5.2 Empirical specification

Because GenAI affects the entire setting, Stage 2 is not estimated as a causal design with an untreated population. For the organizational-type analysis in Table 3, the estimands are comparative: whether post-2022 changes differ across organizational types relative to traditional corporations, and whether those differences remain after allowing each type its own linear time trend. Table 3 reports the primary specification, a post-2022 slope-break model estimated on the balanced pre-period sample of 453 firms and 4,077 observations over 2017-2025. The model includes firm and year fixed effects and allows each organizational type its own linear trajectory. The table reports three relational and three executional dictionary constructions, with standard errors clustered by firm. We estimate the following specification:

Table 3: Stage 2 — Panel Regressions: Differential Language Shifts Across Organizational Types

	Relational language			Executional language		
	(1)	(2)	(3)	(4)	(5)	(6)
	Comprehensive	Client/ advisory	Stakeholder responsibility	Comprehensive	Interface/ scalable	Delivery/ implementation
Professional Services x Post-2022 slope	13.868 (0.005)	9.138 (0.027)	4.730 (0.028)	1.204 (0.542)	-1.083 (0.406)	0.850 (0.491)
Emerging Tech x Post-2022 slope	4.483 (0.408)	1.786 (0.702)	2.697 (0.189)	-1.865 (0.614)	1.605 (0.577)	-2.024 (0.297)
Incumbent Tech x Post-2022 slope	1.765 (0.705)	-0.086 (0.982)	1.851 (0.382)	5.364 (0.130)	2.722 (0.307)	2.030 (0.253)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Type-specific linear trends	Yes	Yes	Yes	Yes	Yes	Yes
Balanced pre-period sample	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.718	0.761	0.705	0.813	0.829	0.748
Observations	4,077	4,077	4,077	4,077	4,077	4,077
Firms	453	453	453	453	453	453

Notes. Coefficients are reported with exact p-values in parentheses; standard errors are clustered by firm. All columns use the 2017-2025 balanced pre-period sample and estimate type-specific changes in slope after 2022. Models include firm fixed effects, year fixed effects, and organization-type-specific linear trends. The omitted organization type is traditional corporations. Dependent variables are dictionary terms per 10,000 words.

$$Language_{it} = \alpha_i + \lambda_t + \sum_g \beta_g (Type_{ig} \times Trend_t) + \sum_g \gamma_g (Type_{ig} \times PostTrend_t) + \varepsilon_{it}$$

where g indexes professional services firms, emerging technology firms, and incumbent technology firms, with traditional corporations omitted; $Trend_t$ is a linear time trend; and $PostTrend_t$ equals zero through 2022, one in 2023, two in 2024, and three in 2025. The γ_g coefficients are the post-2022 slope breaks reported in Table 3.

5.3 Results

Table 3 reports positive post-2022 changes in the relational-language trends of professional services firms relative to traditional corporations. In the balanced slope-break specification, the Professional Services x Post-2022 slope is 13.868 ($p = .005$) using the comprehensive relational dictionary, 9.138 ($p = .027$) using the client/advisory dictionary, and 4.730 ($p = .028$) using the stakeholder-responsibility dictionary. The corresponding emerging-technology and incumbent-technology coefficients are statistically imprecise across all three relational measures.

The executional estimates in Table 3 provide weaker evidence of an aggregate post-2022 break among technology firms. The emerging-technology coefficients are -1.865 ($p = .614$) for the comprehensive executional dictionary, 1.605 ($p = .577$) for interface/scalable-systems language, and -2.024 ($p = .297$) for delivery/implementation language; the incumbent-technology estimates are likewise statistically imprecise (5.364, $p = .130$; 2.722, $p = .307$; 2.030, $p = .253$). Table 4 supplements these organizational-type estimates with subgroup analyses using the full panel.

5.4 Ex-ante heterogeneity

Table 4 uses the full 2017-2025 panel of 626 firms and 5,330 observations. Each column estimates a separate model comparing one focal subgroup with all remaining firms, with firm and year fixed effects, a focal-group linear trend, and a focal-group post-2022 slope interaction. The groups in Table 4 are constructed from the sample's subsector labels before estimating the language models; no post-2022 term counts, coefficients, or observed language changes enter the assignment. Professional services firms are divided into bundled advisory firms (management consulting, investment banking, and creative/digital agencies), accredited firms (law and accounting), and decomposable service providers (HR/staffing, architecture, engineering/environmental services, IT services, and BPO). Technology firms are divided into interface-oriented firms (AI/data, platforms/internet, and software/SaaS), hardware/fintech firms (semiconductors/hardware and fintech/payments), and security/cyber firms. Firms with broad or non-fitting labels remain in the estimation sample but are not reported as focal columns; the Online Appendix reports the full mapping and firm counts.

Relative to the remaining firms, bundled advisory firms exhibit a positive post-2022 slope change in comprehensive relational language (12.849, $p = .043$). The corresponding estimates for accredited firms (2.570, $p = .750$) and decomposable service providers (0.061, $p = .994$) are statistically imprecise. For interface/scalable-systems language, the estimated slope change is positive for interface-oriented technology firms (8.189, $p = .009$) and statistically imprecise for hardware/fintech firms (-1.226, $p = .664$). Security/cyber firms exhibit a negative relative slope change in comprehensive executional language (-11.124, $p = .046$).

Table 4: Stage 2 — Ex Ante Heterogeneity in Website Positioning

	Professional services: relational language			Technology: interface/executional language		
	(1) Bundled	(2) Accredited	(3) Decomposable	(4) Interface tech	(5) Hardware/fintech	(6) Security/cyber
Group x Post-2022 slope	12.849 (0.043)	2.570 (0.750)	0.061 (0.994)	8.189 (0.009)	-1.226 (0.664)	-11.124 (0.046)
Group x Linear trend	-0.291 (0.932)	-0.468 (0.919)	6.401 (0.048)	2.561 (0.107)	0.041 (0.977)	3.239 (0.300)
Group firms	63	73	28	100	44	27
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.732	0.730	0.731	0.831	0.822	0.805
Observations	5,330	5,330	5,330	5,330	5,330	5,330
Firms	626	626	626	626	626	626

Notes. Each column reports a separate OLS regression on the full 2017-2025 panel (626 firms; 5,330 observations), comparing the focal subgroup with all remaining firms. Dependent variables are dictionary term counts per 10,000 words: comprehensive relational in columns (1)-(3), interface/scalable systems in columns (4)-(5), and comprehensive executional in column (6). Models include firm and year fixed effects, a focal-group linear trend (year minus 2020), and a focal-group post-2022 slope interaction, coded zero through 2022 and one, two, and three in 2023, 2024, and 2025. The slope coefficient measures the post-2022 change in the focal group's trend relative to the remaining firms, accounting for the pre-break difference in trends. P-values are in parentheses; standard errors are clustered by firm. Subgroup assignments use subsector labels independently of language outcomes.

The subgroup estimates motivate closer examination of how firms' inherited offerings relate to their public positioning. Bundled advisory firms are the professional services subgroup whose offering most directly combines problem framing, senior judgment, client access, and delivery inside a single engagement; interface-oriented technology firms are the subgroup whose products already take the form of scalable digital interfaces. For interface-oriented technology firms, Appendix Table B13 reports positive relative slope changes in comprehensive executional and interface/scalable-systems language, whereas the delivery/implementation estimate is statistically imprecise. These findings identify measure-specific subgroup patterns alongside statistically imprecise aggregate technology estimates.

5.5 Summary of Stage 2

In the balanced pre-period sample, professional services firms exhibit positive post-2022 changes in relational-language trends relative to traditional corporations across all three dictionary constructions. The aggregate technology slope-break estimates are statistically imprecise; in the full-panel subgroup analyses, interface-oriented firms exhibit positive relative slope changes in comprehensive executional and interface/scalable-systems language. This is consistent with firms foregrounding the value they can plausibly defend from their inherited complements: bundled advisory firms around client access, judgment, trust, and accountability, and interface-oriented technology firms around scalable systems. Website language cannot show boundary change or value capture; the narrower point is that firms publicly reposition around what appears scarce and difficult to replace as generation becomes cheaper. Stage 3 turns to how control of such complements bears on activity allocation.

6. Stage 3: Conditions and Mechanisms of Activity Allocation

Stage 2 documented systematically different shifts in public positioning across organizational types. A firm is a composite, however: a single professional services firm sells advisory judgment, standardized analysis, regulated attestation, and routine administration, and these need not move together. Stage 3 therefore moves beneath the firm to the activity. For each activity informants discussed, we ask what is now produced differently, who performs it, and how it is specified, verified, and paid for across the boundary between buyer and provider.

6.1 Design and Participants

Stage 3 draws on 33 semi-structured interviews with senior executives and directors (late 2025 to early 2026) and two structured roundtables (40 participants; 73 participants in total; Table 5), selected to maximize variation across organizational types while maintaining seniority and direct engagement with AI adoption. Interviews of 30 to 75 minutes and roundtables of five hours followed protocols covering sectoral change, repositioning, and constraints on response (Langley, 1999; Miles and Huberman, 1994), were recorded and transcribed verbatim, and were analyzed following Gioia et al. (2013), with independent coding reconciled through discussion; the data structure and coding protocol are in the Online Appendix. Because the setting is rapidly evolving and the sample purposive, we do not claim conventional theoretical saturation; the design seeks variation and early mechanism identification rather than prevalence estimates.

Table 5: Stage 3 — Roundtables and One-to-One Interviews

ROUNDTABLES (A & B)	ONE-TO-ONE INTERVIEWS
Goldman Sachs - <i>Chief Investment Officer</i>	Accounting & Tax Advisory – <i>Director</i>
Goldman Sachs - <i>Head of AI Accelerator, PE Investments</i>	AI Ethics & Systems Engineering Consultancy – <i>Director</i>
Google Research - <i>Distinguished Scientist</i>	Asset Management Firm – <i>Independent Director</i>
Microsoft – <i>Economist</i>	Brand Strategy Consultancy – <i>Managing Director</i>
Bloomberg - <i>Head of Responsible AI, CTO Office</i>	Construction & Strategy Consultancy – <i>Managing Director</i>
The Estée Lauder Companies - <i>SVP, Marketing, Head of AI</i>	Construction Hardware & Technology Company – <i>Director</i>
Procter and Gamble - <i>Director of Strategy, R&D</i>	Defense & Security Consultancy – <i>Managing Director</i>
ExxonMobil - <i>Product Commercial Analytics Manager</i>	Digital Policy Organization – <i>Director</i>
Vodafone - <i>Head of Data Portfolio</i>	Due Diligence & Compliance Consultancy – <i>Director</i>
SAP - <i>Regional Director, AI Strategy</i>	Education & Assessment Body – <i>Director</i>
Expedia - <i>SVP, Product</i>	Energy Regulation Body – <i>CEO</i>
Edward Jones – <i>Partner</i>	Financial Services Consultancy – <i>Director</i>
AlixPartners - <i>Partner and Managing Director</i>	Financial Services Technology Company – <i>CEO</i>
Seven2 Private Equity - <i>Partner, Value Creation</i>	Financial Technology Consultancy – <i>Director</i>
State of Massachusetts - <i>Head of AI</i>	HR & Mediation Consultancy – <i>Independent Consultant</i>
The Conference Board - <i>Center Leader, Marketing, Comms</i>	Independent Corporate Advisory – <i>CEO</i>
Iterate.ai - <i>Co-Founder and CEO</i>	Infrastructure & Innovation Consultancy – <i>Managing Director</i>
Evidenza - <i>Co-Founder</i>	IT Governance & Government Advisory – <i>Director</i>
BigSpring - <i>Founder and CEO</i>	Management Consultancy – <i>Director</i>
Google Research – <i>Distinguished Scientist</i>	Marketing & Communications Consultancy – <i>Director</i>
Qualcomm – <i>SVP and GM, AI Business Unit</i>	Pharma & Life Sciences Consultancy – <i>Founder & Director</i>
ExxonMobil – <i>Chief Engineer, Data Science</i>	Pharma Advisory – <i>COO</i>
Procter and Gamble – <i>R&D Vice President</i>	Private Equity & Law – <i>Managing Director</i>
Mondelez Digital Services – <i>Senior Vice President</i>	Supply Chain Consultancy – <i>Founder & Director</i>
Tata Communications – <i>VP, Global Head of New Ventures</i>	Professional Governance Body – <i>Director</i>
BCG – <i>Partner</i>	Real Estate & Infrastructure Consultancy – <i>Vice President</i>
BCG X AI Science Institute – <i>VP of AI, Data Science</i>	Sales Training & Advisory – <i>Director</i>
Fiserv – <i>SVP, Global Digital Solutions</i>	Strategy Consultancy – <i>Executive Coach</i>
Fujitsu North America – <i>Regional Chief Technology Officer</i>	Technology & Knowledge Management Consultancy – <i>CEO</i>
GoodRx – <i>Chief Product and Technology Officer</i>	Technology & Telematics Company – <i>CFO</i>
BootstrapLabs – <i>Co-Head, VC, and AI Groups</i>	Technology Company – <i>Director</i>
GrowthCurve Capital – <i>Head of Data, Analytics, ML</i>	Technology Company – <i>Director</i>
GrowthPoint Technology – <i>Managing Director</i>	Education & Technology Company – <i>Director</i>
Pinegrove Venture Partners – <i>Chief Operating Officer</i>	
SLW – <i>Managing Director</i>	
Evolver – <i>Founder and CEO</i>	
ObinAI – <i>Co-Founder and CEO</i>	
Emovid – <i>CEO and Co-Founder</i>	
GBK Collective / Mindspace – <i>Founder and CEO; Partner</i>	
Martin Reeves Advisory – <i>Founder and Principal</i>	

Notes: Roundtables A and B included 19 and 21 participants respectively (40 total); one-to-one interviews were conducted with a separate group of 33 respondents. Firm names and roles are disclosed for roundtable participants; interviewees are listed with anonymized descriptions.

Coding proceeded from the activity outward: what informants said had changed about it, who now performed it, and what they identified as producing the change. Accounts were coded as realized, prospective, reversed, or non-moving; realized accounts required a specific referent, such as a service line no longer offered, an operations center restructured, or a budget not reallocated. Second-hand and vendor-supplied accounts are labeled as such. Claims about reallocation rest on realized accounts with identifiable referents.

Table 6: Stage 3 — Activity-Level Accounts by Configuration and Evidentiary Status


CONFIGURATION 1 — HIGH BYPASS, GOVERNABLE EXCHANGE: CONTESTABLE REALLOCATION		n
Realized reallocation	Expert identification; evidence planning; payer research; literature screening and extraction (pharma advisory). Security operations center monitoring; penetration testing (cybersecurity). Graphic and visual asset production (defense and security consulting); cross-firm invoicing and proof-of-delivery retrieval (construction distribution, accounting); look-back tax recovery (tax advisory).	9
Realized, then reversed	Medical publications content development (pharma advisory); analytics and insight production (cross-sector data provider).	2
Prospective	Bill-of-materials procurement and value engineering (construction supply).	1
CONFIGURATION 2 — HIGH BYPASS, RESISTANT EXCHANGE: PERSISTENT BUNDLES AND SELECTIVE INSOURCING		
Realized: activity moved toward the client	Marketing copy and content drafting (marketing consultancy); market research (management consulting); wellbeing diagnosis and action planning (technology advisory); lower-tier analysis (government advisory); publication support (governance); sales training and capability transfer (sales consultancy); exception handling in accounts payable (financial transformation).	7
Realized: activity absorbed by the provider, or delivery inverted	Bookkeeping and year-end preparation, re-integrated (accounting); consultancy delivery inverted to machine-primary, second-hand (technology and business advisory).	2
Realized non-contracting	Network-level scam interdiction (telecommunications).	1
Prospective, attempted	Board packs and financial analysis (cybersecurity).	1
CONFIGURATION 3 — LOW BYPASS, GOVERNABLE EXCHANGE: FEDERATED PROVISION OF HUMAN EXPERTISE		
Realized reallocation	Specialist design provision coordinated through shared models; technical and production architecture drawings (construction). Independent validation and gap analysis (real estate and infrastructure). Contracted delivery of IT services (IT services). Temporary specialist and freelance provision (management consulting).	5
CONFIGURATION 4 — NEITHER CONDITION: INTEGRATED EXPERT PROVISION		
Realized non-movement, constraint, or failure to contract	Problem framing and question reframing (strategy and management consulting); multi-tax planning where obligations interact (accounting and tax); statutory audit opinion and accountable assurance (accounting); legal advice requiring professional sign-off (law); accountability-bearing sign-off on translated content (large technology firm); audit of machine-produced work and of models (assurance); regulatory approval and building control assessment (construction).	7
Realized: demand rising	Independent second-opinion verification (accounting; strategy and defense consulting).	1

Notes: The table reports all 36 activity-level accounts used in the Stage 3 analysis, grouped by the configuration to which each activity was assigned and by evidentiary status. Assignment was made from task and exchange characteristics independently of the reported allocation outcome; the per-activity assignment audit is in Online Appendix Table C3. Realized accounts require an identifiable referent. Counts are of accounts, not of firms or of incidence, and are not prevalence estimates: a single informant may discuss several activities, and an activity may be qualified by more than one account. The described change, and the qualifying or disconfirming account for each activity, are reported in full in Online Appendix Table C5. Reading across the panels: realized reallocation is concentrated in Configurations 1 and 3, movement in Configuration 2 runs toward the client, back into the provider, or stops short of contracting rather than entering a market, and Configuration 4 contains no realized reallocation at all.

Guided by the first two stages, we then coded what informants identified as producing each change. Stage 1 had indicated where to look: participants separated the generation of an output from the conditions of its use, and attributed unevenness to accountability, regulation, buyer structure, and the terms on which an activity could be transacted. Each account was accordingly coded for the characteristics informants named - among them decomposability and interdependence, ex-ante specifiability, ex-post verifiability, the assignment of warrant or liability, and whether the focal output could be generated without the embodied expertise formerly required - independently of the allocation outcome it reported.

The coded characteristics grouped into two dimensions. The first is *exchange governability*: whether an activity can be coordinated, specified, verified, priced, contracted, and warranted across an organizational boundary, the exchange-side conditions reviewed above (Baldwin, 2008; Jacobides, 2005; Jacobides and Winter, 2005). GenAI sometimes contributed by making an output more demonstrable, lowering search or coordination costs, or enabling a standing interface. In other activities the exchange did not change, or governability changed through digital infrastructure, contracting arrangements, procurement rules, accreditation, or liability regimes that predate GenAI and adjust on their own timescales. The second dimension is *expertise bypass*: the extent to which GenAI can produce the focal output without mobilizing the embodied expertise previously required, often with sharply lower expert labor and elapsed time. Informants described expertise bypass increasing across many activities at once; by itself, it does not establish where an activity will sit or who will capture the gain.

Figure 1: Configurations of the Conditions for Architectural Change in Knowledge Work



Notes: The horizontal axis distinguishes activities by the governability of their exchange, that is, whether the activity can be coordinated, specified, verified, priced, and warranted across an organizational boundary. The vertical axis distinguishes activities by the degree of AI-enabled expertise bypass, that is, whether the output can be produced without mobilizing the embodied expertise previously required. The unit placed in the figure is the activity, not the firm; informants described activities belonging to different configurations within the same organization. Activities were classified from task and exchange characteristics independently of their reported allocation outcomes. Cell entries name the organizational patterns informants described. Within each configuration, control of residual complements shapes, rather than mechanically determines, the destination of the activity. Where accounts distinguished a machine-generable production step from accountable warranting or sign-off, the figure places the focal activity component whose allocation was at issue.

Crossing the two dimensions yields four configurations (Figure 1). Activities were assigned on the coded characteristics, and reported destinations were then examined within, rather than used to define, the configurations; where an account separated a machine-generable production step from an accountable warrant or sign-off, the component that crossed the boundary was coded separately from the warranting activity. Table 6 reports every activity-level account and its evidentiary status; Table 7 classifies the mechanisms informants named into direct production effects, mechanisms bearing on exchange governability, and complementary infrastructure that predates or merely incorporates GenAI. The following sections report the four configurations in turn.

Table 7: Stage 3 — Mechanisms by Class

<i>Mechanism</i>	<i>What it changes</i>	<i>Where it was observed</i>
CLASS A — DIRECT GENAI PRODUCTION EFFECTS (EXPERTISE BYPASS)		
Encoding of practitioner judgment into machine-executable form	The expertise required to produce an output is no longer mobilized at the point of production	Wellbeing action planning; payer research; codified accounting experience in a software product
Inversion of delivery from human-primary to machine-primary	The ratio of machine to human input in producing a deliverable	Security operations center monitoring; a fifteen-consultant advisory firm (second-hand)
Machine-feasible continuous operation	Forms of provision that would be prohibitively costly with predominantly human labor	Continuous penetration testing in place of annual engagements
Client-side self-provision of first-draft output	Who performs the production step, without changing who warrants the result	Marketing copy; first-pass tax queries; consultation responses
Limit drawn by informants	Generation stops at validation: GenAI can produce an analysis but neither frame the question ex ante nor warrant the answer ex post	Problem framing (4 accounts); multi-tax interaction; exception handling
CLASS B — MECHANISMS BEARING ON EXCHANGE GOVERNABILITY		
Platform or subscription re-specification of a project as a standing interface	Converts an episodic, negotiated engagement into a specified, priced, recurring exchange	Expert identification; evidence planning; payer research; security operations
Outcome-contingent contracting	Shifts verification from monitoring the process to measuring the result	Look-back tax recovery; gain-share terms in outsourced renewals; IT services (disputed)
Machine-mediated coordination of independently produced components	Reduces the cost of integrating work performed by separate firms	Clash detection and model merging in shared design software
Unbundling of validation from production	Makes verification separately specifiable and separately purchasable	External expert validation and gap analysis of consultancy deliverables
Third-party assurance of machine-produced output	Supplies the warrant a buyer cannot generate for itself	External-auditor-proof assurance layer; accreditation to audit AI systems
Limit drawn by informants	Governability stops where interdependence is reciprocal, where accountability cannot be assigned, and where the instrumentation to verify machine-produced work does not exist	Multi-tax planning; statutory audit; audit of models; scam interdiction; sign-off on translated content
CLASS C — COMPLEMENTARY INFRASTRUCTURE THAT PREDATES OR MERELY INCORPORATES GENAI		
Structured repositories and data ontologies	Makes internal exchange machine-legible; a precondition for Class B	Bank data dictionaries; data cleansing reported to have risen from about one month to six to nine months
Shared cloud design and modelling platforms	Permits independently produced components to be merged into a single model	Building information modelling in construction
Open banking and ledger interoperability	Removes manual transcription between counterparties	Automated bank feeds and invoice scanning; direct invoice transmission between ledgers
Certification and procurement infrastructure	Determines which providers may transact with which buyers, independently of capability	Information security certification as a condition of onboarding; creditworthiness thresholds
Systems integration and authentication across heterogeneous systems	Determines whether an activity can be executed across organizational boundaries at all	Integration across a state government's case-management systems, declined by major platform vendors

Notes: Evidence from 33 semi-structured interviews and two structured roundtables (40 participants). The three classes separate GenAI's direct production effects from mechanisms bearing on exchange governability, and both from complementary digital infrastructure that predates or merely incorporates GenAI. The classification exists to prevent attributing all standardization and coordination change to GenAI: the Class C mechanisms were in several cases attributed by informants explicitly to something other than GenAI. This is a mechanism inventory, not a prevalence estimate; extended evidence is in Online Appendix Tables C5 and C6.

6.2 Configuration 1: Contestable Reallocation (High Bypass, Governable Exchange)

In the first configuration both conditions were present. The output of an activity could be generated without the expertise that had produced it, and the resulting exchange could be specified and verified well enough to be bought at arm's length. Informants described activities of this kind leaving the bundled engagement for platform-mediated provision. In pharma advisory, expert identification, a bespoke search built on practitioners' familiarity with who mattered in a therapy area, had ceased to be a purchasable engagement, because "people have aggregated all this into databases and that project has gone", so that "there's actually, from [a] strategic perspective, no point in us offering to a client, we'll do expert identification, because they say, well, why would we ask you to do that? We can just buy it" (COO, pharma advisory). Evidence planning and payer research followed the same path. The aggregation predates GenAI; what informants described as newer was the generative layer built on the archive, and the subscription interface that re-specified a project as a standing exchange.

The most fully realized case was in cybersecurity:

"We have multiple LLMs that now essentially work through and do the bulk of the work in the SOC for us. That actually means we've been able to reduce the technical heads that are needed in the SOC by about two thirds. [...] We're taking savings off of the SOC and putting that resource back into R&D, in terms of producing and getting this LLM to build the new product for us." (CFO, cybersecurity technology)

The reduction came through attrition rather than redundancy, and the freed resource went into continuous penetration testing in place of the annual engagement buyers had commissioned. The provider placed the change in the exchange rather than in the work: "you'll actually see quite a change, not in the products itself, but how it's packaged." The same firm indicated why the two dimensions can vary independently, since its two principal activities were verified by buyers in different ways: "the SOC ... that's demos, that's showing it, that's showing proof of concept. For everything else in cybersecurity, which is more purely people-based ... that is more brand and reputation, accreditation, very important." Where an activity's output could be demonstrated, its exchange could be specified and priced as a product; where it could not, buyers fell back on accreditation and the activity stayed inside the bundle. Informants located what was defensible in this configuration narrowly. Four of them discounted proprietary data as leaky, pointing instead to the currency and coverage of an archive within a specific domain, together with an interface that rendered it purchasable.

Reallocation in this configuration was contested rather than settled. A pharmaceutical company acting on consultancy advice moved medical publications content development to a large technology platform and sought a thirty per cent price reduction from its agencies, and the activity returned when the platform "didn't work at all. It was useless" (COO, pharma advisory); a data provider described clients insourcing analytics production and returning it within about six months, having found that "actually, no, that's pretty hard" (Director, enterprise data and analytics provider, Roundtable A). The fluidity that this quadrant entails means that firms will

experiment with several ideas, often overestimating their own capabilities or abilities to structure an exchange effectively.

6.3 Configuration 2: Persistent Bundles, Selective Insourcing (High Bypass, Resistant Exchange)

In the second configuration, only bypass was present. The output could be generated, often at a fraction of its former cost, but its exchange resisted specification because it required client-specific interpretation, reciprocal iteration, or the bearing of accountability. Activities of this kind did not move into a market; informants described transformation inside the bundle and, in several accounts, movement into the client's own organization. Content drafting in a marketing consultancy had already crossed part of that boundary: two clients produced first drafts themselves, so that "one of them gets it to write her newsletter for her and then she sends it to me. But I still adjust it because she hasn't taught it quite the right tone of voice", and the provider's work had shifted from production to interpretation: "I'm being employed more for my professional skills as a marketer now, rather than predominantly my skills as a copywriter." The residual activity was itself described as provisional:

"There probably will come a time where they either say, you can have access to my account, you can run the AI stuff on my account, or, we'll do that bit and then we'll come back to you and ask you to support what we've created ... because actually the bit we value is your skills in teaching it and getting it right for us as a business. Which obviously means that there'll come a time when they say we don't need you to do that bit at all." (Director, marketing and communications consultancy)

The account describes a boundary that had already moved once and was expected to move again. The current complement lay in training the client's tools rather than producing the content. Across sectors, informants placed the limit of bypass at problem framing: "I did strategy for 36 years for [a major global consultancy]. I cannot remember a single time when the client question was posed correctly ... there is no LLM that will give you the right answer to the wrong question" (Former partner, global strategy consultancy, Roundtable B). Multi-tax planning supplied the structural version of the limit, since United Kingdom obligations interact and machine advice was "answering questions on specifics rather than linking things together. It's that linking of the different taxes together" (Director, accounting and tax advisory).

Within this configuration, the destination of an activity depended on who controlled the residual complement (Teece, 1986). When the provider controlled judgment, legitimation, or liability, the bundle persisted beneath a human interpretive layer. When the client controlled proprietary data, organizational context, or verification requirements, informants described movement inward. A wellbeing advisory firm encoded its expert panel's judgment into a product, so that where clients had previously "had to engage a third party to deliver some solutions ... now, the HR team, generally, are empowered to then take some actionable steps" (Director, financial technology consultancy), an arrangement acceptable to the buyer because "it's less important that it's 100 per cent accurate." Provider-side accounts pointed the same way. One consultant now advised clients "to insource the capability, because insourcing that capability comes now

at a lower cost - they don't need to employ a specialist consultant, they can employ fairly junior analysts", while market research had fallen from "50k a pop" to "five, or 10" (Managing director, infrastructure and innovation consultancy).

Movement in this configuration ran in both directions. Bookkeeping had sat outside a small accountancy practice for years, because the practice's cost of performing it exceeded what clients would pay; bank-feed automation and invoice scanning altered that calculation and the practice absorbed it, since "whereas three years ago, we didn't really do bookkeeping. We are taking on a lot more bookkeeping clients because it's cost-effective for us to do it because of the ease with software" (Director, accounting and tax advisory). The mechanism there was complementary infrastructure rather than GenAI, and the direction was re-integration.

6.4 Configuration 3: Federated Expert Provision (Low Bypass, Governable Exchange)

The third configuration reverses that combination. The expertise an activity required remained tacit and human, and no informant suggested that a machine could supply it; what changed was the exchange, which became specifiable and verifiable enough to be conducted between firms rather than inside one (Choudary, 2025). In construction, projects are decomposed into components produced by different specialists and merged into a master model, and the merging had been "a lot of manual checking." Tools embedded in shared design software now performed clash detection, error detection and model reconciliation across independently produced components, and the coordination overhead that had sustained integration fell far enough that assembled teams competed with integrated firms:

"What then happens is instead of the big guy offering a one-stop shop package to a client, [...] I can offer almost the same performance and delivery as the one-stop shop because my entire team, even though it's made up of a number of sub-consultants, we're still working on it as if it's a project in-house because everything's being done using cloud-based tools, and it's all integrated into a single model." (Managing director, construction)

Such teams "form, they do a project, and then they dissolve ... That's always happened in the industry, but that's become much more efficient now." The account involves no bypass of design expertise; the change lay in the cost of combining independently produced work. Validation supplied a second instance. A buyer of work from several large consultancies described contracting an external subject-matter expert to check machine-assisted output produced by junior teams, since "they'll create the goods, and they'll give them to me to validate ... And do a gap analysis" (Vice president, real estate and infrastructure consultancy). The specialist expertise was untouched; validation had become separately specifiable, and so separately purchasable, from production.

Outcome-contingent contracting extended the same logic, and was the one mechanism on which informants openly disagreed. An information technology services executive described a transition already underway, since "this whole model of the IT services industry that was linearly based in terms of people has completely turned on its head now ... You've got to price yourself based on outcomes" (Chief technology officer, IT services, Roundtable B). Others at

the same roundtable held that outcome terms worked only "for purely standardizable problems", since "anything you have a standard metric for, you can probably get more cheaply elsewhere", and that "outcome so far is a sales trick to me". A contrasting account, reported in Configuration 1 because machine review supplied the production step, concerned look-back tax recovery: a model reviews every transaction for over-paid state tax and the provider takes about "40 cents on the dollar" of what it recovers. Across the accounts, contingent pricing appeared where the outcome was countable.

One reallocation in this configuration involved no bypass at all. In the same construction interview, technical architectural drawing had been separated out and made a distinct traded activity because of how liability is assigned, since "there's a lot of creative architects now who don't want to do that piece because it carries liability with it ... [my architects] have set up a division of their company that just does these technical drawings ... So they take work from other concept architects who don't want the responsibility legally." Asked what had driven the change, the informant was explicit: "that has nothing to do with AI. That has everything to do with contract law and who's responsible." Governability, in these accounts, was a property of how an activity's accountability could be assigned and specified, and it moved on its own timescale.

6.5 Configuration 4: Integrated Expert Provision (Low Bypass, Resistant Exchange)

In the fourth configuration neither condition held for the focal activity, and informants described no realized reallocation. The activities that sat here shared a feature that is easily lost at the level of the firm: their output is a warrant. Statutory audit opinions, professional sign-off, regulatory attestation, problem framing, multi-tax planning, and independent second opinion deliver less a discrete analysis than an accountable party's undertaking that an analysis may be relied upon. We label the configuration integrated expert provision, rather than adopting informants' relational vocabulary, because what held these activities in place was identifiable. An audit is "a legally enforceable service that a company of a certain size has to contract with an accountant to produce" (Managing director, brand strategy consultancy), three informants named professional indemnity insurance specifically, and a former lawyer expected regulated professionals to under-adopt for that reason:

"Probably the sectors where you've got professional requirements, you've got professional governance - solicitors, accountants, auditors - where they've always had a senior partner responsible ... I kind of feel like they won't get sucked into this quite so much because their personal and professional reputations rely on it and their insurance relies on it." (Director, financial services governance)

Statutory audit illustrates why we separate machine-generable production from the accountable warrant. A founder building a domain-specific model for audit reported machine execution of audit procedures within reach, since "the job of audit can be autonomously done. Can be done, right? ... But I have not seen any evidence of any audit firm, or any entrepreneurial company ... actively transforming ... the sector that I'm working [in], which is audit." Asked what stood in the

way, he named the exchange: "our challenge with my startup is how to provide the evidence and the assurance to a particular customer that an AI doing it autonomously can be reliable without a human" (Director, financial services consultancy, and founder of an AI audit venture). Two further accounts indicated why that assurance was hard to supply. A public-sector head of artificial intelligence reported an internal analysis finding that most models had "very poor to no observability capabilities ... [so] even if you wanted to audit the AI models ... you couldn't do that. Not today", leaving audit "largely human today ... not being productized" (Roundtable A); an investor described suitably accredited AI-assurance capacity as extremely scarce.

The configuration was contested even so. One entrant argued that incumbent assurance rests on sampling, so that "over enough time the notion of trust is going to start eroding ... [because] are you really looking at every single transaction? ... Or are you doing sampling?" (Founder, AI-native finance and audit platform, Roundtable B), while a participant formerly responsible for artificial intelligence products at a large technology firm recalled a ten-million-dollar budget he could not reallocate to machine translation over six or seven years, because the buyer "needed somebody when things go wrong to call on and blame" (Co-founder, AI applications company, Roundtable B). Bypass elsewhere had, if anything, raised demand for the activities in this configuration, since clients now arrived "with partial advice already ... they just want to check that what they've seen on the internet is correct" (Director, accounting and tax advisory).

Legal advice illustrates why internal adoption and boundary movement are separate questions. Informants described adoption inside law firms as fast, with precedent research and paralegal work substantially machine-performed, and attributed the speed to organizational form, since "most law firms don't have a CTO as such" (Managing director, private equity and law). The allocation of accountable legal advice and sign-off was unchanged. Two mechanisms were offered for that stability. Regulated professionals treated regulation as protection, a characterization that drew "it's a moat, really" when we put it to a policy director; and a participant with a decade in large-firm consulting pointed instead to who can fund a change in the exchange, since "the auditors are partner-based firms, and anybody that actually has decision-making powers is two years away from retirement. They're not going to invest in a brand new platform ... because they're never going to capitalize on it. Same on the legal side" (Commercial director, enterprise AI platform, Roundtable B).

6.6 Corollary: Entrepreneurial Positions Opened by the Configurations

The shifts informants described were not confined to reallocation among incumbents. They also identified emerging or prospective organizational forms associated with the positions opened by the configurations (Table 8). Where bypass was high and the exchange sufficiently governable (Configuration 1), entry could take the form of AI-native micro-firms, with very small teams competing against much larger incumbents, although the same conditions also produced countervailing scale effects and, in cybersecurity, rising barriers to entry into activities requiring substantial infrastructure. Where an activity admitted bypass but its exchange resisted standardization (Configuration 2), an entrant could not simply buy the activity in a market and instead acquired the client relationship, in the acquire-and-automate form: "I'm just going to buy

your book of business ... I've got a platform, I've got an AI, I can do what you do with five people, you're doing it with 200" (Partner, alternative asset manager, Roundtable B).

Table 8: Stage 3 — Representative Evidence on Emerging and Prospective Organizational Forms

Organizational form	Configuration	Why the configuration opens the position	Representative evidence
AI-native micro-firms	1 - High bypass, governable exchange	Where the output can be generated and the exchange can be specified, very small teams can supply it without the overhead of a large incumbent. Countervailing: infrastructure cost favors scale, so democratization and consolidation appear together.	"You could have a team of two people or three people that are now able to compete with a company that [has more than one hundred people]. I think it is a democratizer." – Director, financial technology "AI is a very expensive technology to have at that level, in the backend infrastructure needed to maintain AI ... consolidation is the dominant dynamic." – Director, technology
Acquire-and-automate consolidators	2 - High bypass, resistant exchange	Where the output can be generated but the exchange resists specification, the activity cannot be bought in a market. The entrant therefore acquires the client relationship and restructures delivery behind it, and the profitability gain funds further acquisition.	"I'm just going to buy your book of business ... I've got a platform, I've got an AI, I can do what you do with five people, you're doing it with 200." – Partner, alternative asset manager (Roundtable B) "They're keeping the relational work and automating all the paperwork. The EBITDA margins have expanded from about eight to ten percent to forty percent." – Advisor, venture capital
Federated specialist teams	3 - Low bypass, governable exchange	Where expertise remains human but coordination becomes cheap, independent specialists can assemble for a project and dissolve on completion, competing with integrated firms on quality rather than on automation.	"I can offer almost the same performance and delivery as the one-stop shop because my entire team, even though it's made up of a number of sub-consultants ... everything's being done using cloud-based tools." – Managing director, construction "A lesser permanent staff and work with freelancers ... because you have a knowledge base, your structured data, your in-house AI." – Director, management consulting
Assurance and integration providers	4 - Low bypass, resistant exchange; cross-cutting governability services	Where governability rather than production capability alone is the binding constraint, entry occurs in the business of relieving it. What these entrants supply is not knowledge work itself but the specification, verification and integration that lets it cross a boundary — a position available because the largest platform owners have declined it.	"A mushroom of one-man bands who are positioning themselves as AI governance experts." – Director, financial services consultancy "There will be companies [whose] entire job is [that] we're going to build agents that can go around and just chisel a hole in the connector." – Roundtable A "Oh, that's your job, go call [a systems integrator]." – Head of AI, US state government (Roundtable A)

Notes: The examples are realized, emerging, or prospective, and the quotations vary in evidentiary status; they are not prevalence estimates. Each form is located in the configuration whose conditions open the position, which is the sense in which entry is a corollary of the framework rather than a separate account. Entry faces the same governability constraint as reallocation: procurement thresholds excluded providers able to perform the work but unable to transact ("if you haven't got 27001, you're going to struggle" – Director, sales training and advisory). Roundtable A refers to the first roundtable, Roundtable B to the second.

Where the exchange became governable while the expertise remained tacit and human (Configuration 3), entry took the form of federated specialist teams, as in construction. And

where governability rather than bypass was the binding constraint, entry took the form of supplying it. Informants described a proliferation of assurance providers, including "a mushroom of one-man bands who are positioning themselves as AI governance experts" (Director, financial services consultancy), alongside the few firms accredited to audit AI systems; and, on the coordination side, ventures built to supply the connectors that platform owners decline to build, since "there will be companies [whose] entire job is [that] we're going to build agents that can go around and just chisel a hole in the connector" (Roundtable A). These entrants supplied the specification, verification, and integration through which knowledge work crosses a boundary, rather than the work itself - a position available because the largest platform owners had declined it: asked to integrate across a state government's case-management systems, the model provider answered "oh, that's your job, go call [a systems integrator]", and "[the major cloud and enterprise software vendors] ... can't do the integration work" (Head of artificial intelligence, US state government, Roundtable A). Entry met its limit at the same governability constraint, since procurement thresholds excluded providers able to perform the work but unable to transact: "if you haven't got 27001, you're going to struggle" (Director, sales training and advisory).

6.7 Cross-Stage Synthesis: Residual Complements, Shifting Bottlenecks, and Organizational Response

Read together, the stages examine different objects rather than measuring the same thing three times. Stage 1 supplied practitioner categories about what appeared exposed or defensible. Stage 2 traced public positioning: professional services firms exhibited positive relative changes in relational-language trends, and interface-oriented technology firms exhibited positive relative changes in executional and interface/scalable-systems trends. Stage 3 developed an activity-level explanation and identified a limited set of realized, reversed, prospective, and non-moving accounts.

Across the stages, this connects two mechanisms at the core of IA: how firms shape the transactional environment and how control of scarce bottlenecks conditions value appropriation (Baldwin, 2008; Jacobides, Knudsen, and Augier, 2006). Expertise bypass and exchange governability define which bundles remain intact and which architectural positions become feasible; control of the residual bottleneck helps explain where activities settle and where value can be defended (Baldwin, 2024; Teece, 1986). GenAI changes the conditions under which activities can be separated and exchanged and, where it makes production cheaper, also changes which assets, capabilities, relationships, and institutional positions remain scarce. Organizational types therefore appear to respond through what they already control: an archive, an accredited signature, a client relationship, proprietary context, or an interface. Where incumbents lacked the relevant complement, the entrepreneurial forms described above acquired it, assembled it across firms, or supplied the conditions that made exchange possible. GenAI thus makes visible how shaping the transactional architecture and retaining control of newly salient bottlenecks become part of the same response.

7. Discussion

Taken together, the stages suggest that GenAI is beginning to place architectural pressure on knowledge work. The pressure operates through the two dimensions identified here: GenAI changes how particular outputs are produced across many activities at once, while the conditions for reorganizing those activities across firms change more selectively. We now draw out the implications for research on AI and work, for IA, and for emerging organizational forms.

7.1 From Task-Level Effects to the Architecture of Knowledge Work

At the task or worker level, the empirical literature on GenAI has made rapid progress on productivity, exposure, and labor-demand effects. Those estimates are the natural starting point for assessing the technology, but they do not specify the organizational response. The same gain may be absorbed within an incumbent, move work across a firm boundary in either direction, or support entry. Our accounts show the gap in both directions. Legal research and audit analytics are highly exposed, yet informants described accountable advice and statutory audit opinions as moving late because liability, accreditation, and professional sign-off resist separation. Specialist design provision in construction is less exposed because the expertise remains human, yet shared digital interfaces had already enabled its redistribution across federated specialist teams. Task exposure tells us what a machine may produce. Expertise bypass, exchange governability, and the location of residual complements together shape where the activity sits and who captures the gain.

7.2 Extending IA: When Production and Exchange Change Separately

We extend IA by specifying how AI-enabled expertise bypass interacts with exchange conditions to shape architectural change. This preserves IA's core insight that productive capabilities and exchange conditions jointly shape firm boundaries, while identifying a production mechanism through which an actor can generate expert output without acquiring or mobilizing the expertise previously required. Because this mechanism can advance without equivalent changes in exchange governability, transformation within persistent bundles can coexist with reallocation across firms. Crossing the two dimensions yields the four configurations of Stage 3: contestable reallocation where expertise bypass is high and the exchange governable; persistent bundles with selective insourcing where output can be generated but the exchange resists; federated expert provision where exchange becomes governable around still-human expertise; and integrated expert provision where neither condition is present for the accountable activity. Within each, residual complements shape where the activity settles. Where production remains stable, the framework reproduces the exchange-side predictions of prior IA work; the off-diagonal cells show what changes when production and exchange diverge.

Three implications follow. First, the boundary at stake becomes finer-grained. Because GenAI changes the production of analytical output, potential "thin crossing points" can emerge inside

previously indivisible engagements, between analysis and the judgment that frames and warrants it (Baldwin, 2008). Second, the two dimensions can change on different timescales, emphasizing the role of clockspeed, in institutions as much as in technology, in shaping industry dynamics (Fine, 1998; MacDuffie, Jacobides & Tae, 2026). Bypass advances with model capability; governability depends on verification technologies, liability regimes, accreditation, and contracting arrangements that change more slowly. Off-diagonal configurations may therefore persist rather than simply mark a transition, helping explain why highly exposed activities can move late. Third, the framework connects architectural change to value appropriation. It shifts attention from adoption alone to the control of interfaces and residual complements. Where generation becomes cheap, the bottleneck can migrate to an archive, proprietary context, client access, an accredited signature, or the interface through which output is verified and sold. The responses we observe are therefore not only choices among alternative boundaries; they are consistent with firms shaping the transactional environment around complements that remain scarce and difficult to replace (Baldwin, 2024; Jacobides, Knudsen, and Augier, 2006; Jacobides and Tae, 2015; Teece, 1986). Technology firms are keen to standardize components; professional services firms strive to keep validation and accountability integrated; clients aim to invest in the interfaces and assurances that make insourcing possible. A common technical shock can therefore generate divergent strategies within as well as across organizational types (Cetorelli, Jacobides, and Stern, 2021; Jacobides and Winter, 2005). Because these complements may themselves become standardized, bypassed, or newly contestable, the locus of scarcity - and therefore of architectural advantage - may continue to move as the technology and its transactional environment co-evolve (Choudary, 2025; Hannah and Eisenhardt, 2018). Our framework connects three interrelated yet distinct questions, helping map the impact of AI on knowledge work: expertise bypass changes who can produce; exchange governability shapes who can transact; and control of scarce complements shapes who can capture value.

7.3 Implications for Emerging Organizational Forms

The conditions we describe open distinct entry positions, consistent with the expectation that architectural change creates entrepreneurial space at newly separable interfaces (Baldwin, 2008; Eisenhardt and Zhang, 2026; Jacobides, 2005; Jacobides, Knudsen, and Augier, 2006; von Krogh, Weiser, and Tinguely, 2026). Informants identified AI-native micro-firms at newly separable analytical interfaces, acquire-and-automate consolidators where ownership substitutes for a resistant exchange, federated specialist teams where standardized coordination supports still-human expertise, and providers of specification, verification, and integration where governability itself is the binding constraint. Whether these forms stabilize, proliferate, or are absorbed remains open. They may coexist with concentration, as in the AI ecosystem itself (Jacobides, Brusoni, and Candelon, 2021), and some may evolve into integrators and legitimators of AI-mediated solutions within client contexts. We present these forms as implications and early signals rather than established population-level outcomes.

7.4 Limitations and Future Research

Several limitations warrant emphasis. We observe an early-stage process whose signals may or may not consolidate. Stage 2 relies on website language, a measure of public claims; distinguishing repositioning from reallocation, and realized from prospective accounts, limits but does not remove that problem. The estimates are correlational, with overlapping period effects, including COVID-19, only partly absorbed by fixed effects, and the balanced panel may introduce survivor bias. Although the comprehensive relational estimate remains positive in the full panel, it is smaller and less precisely estimated. Stage 3 uses a purposive senior-executive sample and does not estimate prevalence. Future research using operational data, employment data, and cross-sectional variation in exposure would strengthen causal interpretation and connect task productivity more directly to organizational allocation.

The dynamics are also subject to boundary conditions. Where advisory judgment and analytical delivery remain reciprocally interdependent, the conditions for separation may not be met (Henderson and Clark, 1990; Thompson, 1967), and complementary assets - regulatory standing, accreditation, and reputational capital - may insulate integrated advisory functions even as the output becomes easier to generate (Helfat and Campo-Rembado, 2016; Teece, 1986). The subgroup analysis yields a positive relative relational slope change for bundled advisory firms and an imprecise estimate for accredited firms, although these estimates do not establish a difference between the two subgroups or distinguish insulation from delay. Industries also oscillate between integration and disintegration as new interdependencies emerge (Fine, 1998). Future research might connect these conditions to work on incumbent adaptation to technological discontinuities (Christensen and Bower, 1996; Eggert and Park, 2018; Tripsas and Gavetti, 2000; Tushman and Anderson, 1986), which asks which firms survive within a given division of labor; our configurations identify when the division of labor itself is at stake.

Three further limitations follow from the framework. Attribution is imperfect: some changes in governability run through broader digitization with which GenAI operates jointly, which is why Table 7 separates direct production effects from complementary infrastructure. The classifications are matters of degree, and activities may shift across configurations as capabilities, regulation, and verification technologies evolve. Informant accounts blend observation with expectation, and some realized accounts are second-hand; Table 6 makes these differences visible without eliminating informant inference. Even with these limits, the early movements show some rhyme and reason. The question is not only what GenAI does to a task, but where that task - and the activity built around it - settles as the architecture of knowledge work begins to shift.

References

- Abbott A. 1988. *The System of Professions: An Essay on the Division of Expert Labor*. University of Chicago Press: Chicago, IL.
- Anthony C. 2021. When Knowledge Work and Analytical Technologies Collide: The Practices and Consequences of Black Boxing Algorithmic Technologies. *Administrative Science Quarterly* 66(4): 1173-1212.
- Autor DH. 2015. Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives* 29(3): 3-30.
- Autor DH, Levy F, Murnane RJ. 2003. The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics* 118(4): 1279-1333.
- Baldwin CY. 2008. Where do transactions come from? Modularity, transactions, and the boundaries of firms. *Industrial and Corporate Change* 17(1): 155-195.
- Baldwin CY. 2024. *Design Rules, Volume 2: How Technology Shapes Organizations*. MIT Press: Cambridge, MA.
- Baldwin CY, Clark KB. 2000. *Design Rules, Volume 1: The Power of Modularity*. MIT Press: Cambridge, MA.
- Bamberger PA. 2018. AMD - Clarifying What We Are about and Where We Are Going. *Academy of Management Discoveries* 4(1): 1-10.
- Behfar K, Okhuysen GA. 2018. Perspective - Discovery Within Validation Logic: Deliberately Surfacing, Complementing, and Substituting Abductive Reasoning in Hypothetico-Deductive Inquiry. *Organization Science* 29(2): 323-340.
- Boussiou L, Lane JN, Zhang M, Jacimovic V, Lakhani KR. 2024. The Crowdless Future? Generative AI and Creative Problem-Solving. *Organization Science*.
- Bresnahan TF, Trajtenberg M. 1995. General purpose technologies: Engines of growth? *Journal of Econometrics* 65(1): 83-108.
- Brynjolfsson E, Li D, Raymond LR. 2025. Generative AI at Work. *The Quarterly Journal of Economics* 140(2): 889-942.
- Brynjolfsson E, Mitchell T. 2017. What can machine learning do? Workforce implications. *Science* 358(6370): 1530-1534.
- Cetorelli N, Jacobides MG, Stern S. 2021. Mapping a sector's scope transformation and the value of following the evolving core. *Strategic Management Journal* 42(12): 2294-2327.
- Chatterji A, Rock D, Talamàs E. 2026. The Human-Machine Knowledge Spiral. arXiv preprint 2606.29227.
- Choudary SP. 2025. *Reshuffle: Who Wins When AI Restacks the Knowledge Economy*. Independently published.
- Choudhury P, Moeen M, Wormald A. 2026. Cultivating and embracing plurality in abductive inquiries within strategy research. *Strategic Management Journal*: 1-22. <https://doi.org/10.1002/smj.70120>.
- Choudhury P, Starr E, Agarwal R. 2020. Machine learning and human capital complementarities: Experimental evidence on bias mitigation. *Strategic Management Journal* 41: 1381-1411.
- Christensen CM, Bower JL. 1996. Customer power, strategic investment, and the failure of leading firms. *Strategic Management Journal* 17(3): 197-218.

- Coase RH. 1937. The Nature of the Firm. *Economica* 4(16): 386-405.
- Cowan R, David PA, Foray D. 2000. The explicit economics of knowledge codification and tacitness. *Industrial and Corporate Change* 9(2): 211-253.
- Creswell JW, Plano Clark VL. 2018. *Designing and Conducting Mixed Methods Research*, 3rd ed. Sage: Los Angeles, CA.
- Darby MR, Karni E. 1973. Free Competition and the Optimal Amount of Fraud. *Journal of Law and Economics* 16(1): 67-88.
- Dell'Acqua F, McFowland E III, Mollick ER, Lifshitz-Assaf H, Kellogg KC, Rajendran S, Kraye L, Candelon F, Lakhani KR. 2026. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. *Organization Science* 37(2): 403-423.
- Demirci O, Hannane J, Zhu X. 2025. Who Is AI Replacing? The Impact of Generative AI on Online Freelancing Platforms. *Management Science* 71(10).
- Dulleck U, Kerschbamer R. 2006. On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods. *Journal of Economic Literature* 44(1): 5-42.
- 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.
- Eisenhardt KM. 1989. Building Theories from Case Study Research. *Academy of Management Review* 14(4): 532-550.
- Eisenhardt, K. M., & Zhang, X. (2026). New wine, old bottles? Effective strategies of new GenAI firms. *Strategic Entrepreneurship Journal*, 20(4), xxx-xxx.
- Eloundou T, Manning S, Mishkin P, Rock D. 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science* 384(6702): 1306-1308.
- Fine CH. 1998. *Clockspeed: Winning Industry Control in the Age of Temporary Advantage*. Perseus Books: Reading, MA.
- Gambacorta L, Qiu H, Shan S, Rees DM. 2024. Generative AI and labour productivity: a field experiment on coding. *BIS Working Paper* 1208.
- Garicano L. 2000. Hierarchies and the Organization of Knowledge in Production. *Journal of Political Economy* 108(5): 874-904.
- Garicano L, Li J, Wu Y. 2026. Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries. CEPR Discussion Paper 21453, Centre for Economic Policy Research.
- Garicano L, Rossi-Hansberg E. 2006. Organization and Inequality in a Knowledge Economy. *The Quarterly Journal of Economics* 121(4): 1383-1435.
- Gennai, J., Katila, R., Kumar, M., & Sinapuelas, I. (2026). Everything everywhere all-at-once? AI and entrepreneurial strategy trade-offs in the age of abundance. *Strategic Entrepreneurship Journal*, 20(4), xxx-xxx.
- 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.
- Girotra K, Meincke L, Terwiesch C, Ulrich KT. 2023. Ideas are Dimes a Dozen: Large Language Models for Idea Generation in Innovation. SSRN working paper.

- Goldfarb A, Tucker C. 2019. Digital Economics. *Journal of Economic Literature* 57(1): 3-43.
- Guzman J, Li A. 2023. Measuring Founding Strategy. *Management Science* 69(1): 101-118.
- Handa K, Tamkin A, McCain M, Huang S, Durmus E, Heck S, Mueller J, Hong J, Ritchie S, Belonax T, Troy KK, Amodei D, Kaplan J, Clark J, Ganguli D. 2025. Which Economic Tasks are Performed with AI? Evidence from Millions of Claude Conversations. *arXiv*.
- Hannah DP, Eisenhardt KM. 2018. How firms navigate cooperation and competition in nascent ecosystems. *Strategic Management Journal* 39(12): 3163-3192.
- Hassan TA, Hollander S, Kalyani A, van Lent L, Schwedeler M, Tahoun A. 2025. Text as Data in Economic Analysis. *Journal of Economic Perspectives* 39(3): 193-220.
- Helfat CE, Campo-Rembado MA. 2016. Integrative Capabilities, Vertical Integration, and Innovation Over Successive Technology Lifecycles. *Organization Science* 27(2): 249-264.
- Henderson RM, Clark KB. 1990. Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms. *Administrative Science Quarterly* 35(1): 9-30.
- Hoffmann M, Boysel S, Nagle F, Peng S, Xu K. 2025. Generative AI and the Nature of Work. SSRN working paper.
- Jacobides MG. 2005. Industry Change Through Vertical Disintegration: How and Why Markets Emerged in Mortgage Banking. *Academy of Management Journal* 48(3): 465-498.
- Jacobides MG, Brusoni S, Candelon F. 2021. The Evolutionary Dynamics of the Artificial Intelligence Ecosystem. *Strategy Science* 6(4): 412-435.
- Jacobides MG, Knudsen T, Augier M. 2006. Benefiting from innovation: Value creation, value appropriation and the role of industry architectures. *Research Policy* 35(8): 1200-1221.
- Jacobides MG, Tae CJ. 2015. Kingpins, Bottlenecks, and Value Dynamics Along a Sector. *Organization Science* 26(3): 889-907.
- Jacobides MG, Winter SG. 2005. The co-evolution of capabilities and transaction costs: explaining the institutional structure of production. *Strategic Management Journal* 26(5): 395-413.
- Jacobides MG, Winter SG. 2012. Capabilities: Structure, Agency, and Evolution. *Organization Science* 23(5): 1365-1381.
- Kapoor, R., & Raj, M. (2026). Value creation and capture in artificial intelligence (AI)-enabled ecosystems. *Strategic Entrepreneurship Journal*, 20(4), xxx-xxx.
- Kogut B, Zander U. 1992. Knowledge of the Firm, Combinative Capabilities, and the Replication of Technology. *Organization Science* 3(3): 383-397.
- Krakovski S, Luger J, Raisch S. 2023. Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal* 44(6): 1425-1452.
- Langley A. 1999. Strategies for Theorizing from Process Data. *Academy of Management Review* 24(4): 691-710.
- Lebovitz S, Levina N, Lifshitz-Assaf H. 2021. Is AI Ground Truth Really True? The Dangers of Training and Evaluating AI Tools Based on Experts' Know-What. *MIS Quarterly* 45(3): 1501-1526.
- MacDuffie, JP, Jacobides, MG, Tae CJ, 2026. From Creative Destruction to Creative Entanglement: Automotive CASE and the Reorganization of Innovation, Working Paper (mimeo) July.

- Marchant GE. 2011. The Growing Gap Between Emerging Technologies and the Law. In *The Growing Gap Between Emerging Technologies and Legal-Ethical Oversight: The Pacing Problem*, Marchant GE, Allenby BR, Herkert JR (eds). Springer: Dordrecht; 19–33.
- Massenkoff M, McCrory P. 2026. Labor market impacts of AI: A new measure and early evidence. *Anthropic*.
- Miles MB, Huberman AM. 1994. *Qualitative Data Analysis: An Expanded Sourcebook*, 2nd ed. Sage: Thousand Oaks, CA.
- Nelson RR, Winter SG. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press: Cambridge, MA.
- 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. The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. *arXiv*.
- Pillai S, Gambardella A, Goldfarb B, King A. 2020. Abduction and the Problem of Null. *Academy of Management Proceedings* 2020(1): 15012.
- Polanyi M. 1966. *The Tacit Dimension*. Doubleday: Garden City, NY.
- Sharma A. 1997. Professional as Agent: Knowledge Asymmetry in Agency Exchange. *Academy of Management Review* 22(3): 758-798.
- Smith A. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. W. Strahan and T. Cadell: London.
- Stigler GJ. 1951. The Division of Labor is Limited by the Extent of the Market. *Journal of Political Economy* 59(3): 185-193.
- Stone PJ, Dunphy DC, Smith MS, Ogilvie DM. 1966. *The General Inquirer: A Computer Approach to Content Analysis*. MIT Press: Cambridge, MA.
- Susskind R, Susskind D. 2015. *The Future of the Professions: How Technology Will Transform the Work of Human Experts*. Oxford University Press: Oxford.
- Teece DJ. 1986. Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy* 15(6): 285-305.
- Teece, D. J. (2026). The entrepreneur as ecosystem architect: Economic foundations for strategic entrepreneurship in the age of AI. *Strategic Entrepreneurship Journal*, 20(4), xxx–xxx.
- Teutloff O, Einsiedler J, Kassi O, Braesemann F, Mishkin P, del Rio-Chanona RM. 2025. Winners and losers of generative AI: Early evidence of shifts in freelancer demand. *Journal of Economic Behavior & Organization* 235: 106845.
- Thompson JD. 1967. *Organizations in Action: Social Science Bases of Administrative Theory*. McGraw-Hill: New York.
- Timmermans S, Tavory I. 2012. Theory Construction in Qualitative Research: From Grounded Theory to Abductive Analysis. *Sociological Theory* 30(3): 167-186.
- Tripsas M, Gavetti G. 2000. Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management Journal* 21(10-11): 1147-1161.

- Tushman ML, Anderson P. 1986. Technological Discontinuities and Organizational Environments. *Administrative Science Quarterly* 31(3): 439-465.
- von Krogh, G., Weiser, A.-K., & Tinguely, P. (2026). "Borrowing intelligence to build a business": The promise of open-source AI for entrepreneurship. *Strategic Entrepreneurship Journal*, 20(4), xxx-xxx.
- von Nordenflycht A. 2010. What Is a Professional Service Firm? Toward a Theory and Taxonomy of Knowledge-Intensive Firms. *Academy of Management Review* 35(1): 155-174.
- Wiles E, Kraymer L, Abbadi M, Awasthi U, Kennedy R, Mishkin P, Sack D, Candelon F. 2024. GenAI as an Exoskeleton: Experimental Evidence on Knowledge Workers Using GenAI on New Skills. SSRN working paper.
- Williamson OE. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press: New York.
- Williamson OE. 1985. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. Free Press: New York.
- Williamson OE. 2000. The New Institutional Economics: Taking Stock, Looking Ahead. *Journal of Economic Literature* 38(3): 595-613.
- Yiu S, Seamans R, Raj M, Liu T. 2025. Strategic Responses to Technological Change: Evidence from Online Labor Markets. SSRN working paper.
- Zander U, Kogut B. 1995. Knowledge and the Speed of the Transfer and Imitation of Organizational Capabilities: An Empirical Test. *Organization Science* 6(1): 76-92.
- Zhang, H., Wuebker, R., & Barney, J. B. (2026). Generative AI and the ability of entrepreneurial opportunities to generate sustained competitive advantages. *Strategic Entrepreneurship Journal*, 20(4), xxx-xxx.
- Zuboff S. 1988. *In the Age of the Smart Machine: The Future of Work and Power*. Basic Books: New York.

Pg **Appendix:**

49 **Appendix A: Stage 1: Executive Roundtable**

49 A1. Design, facilitation, and participant composition

49 A2. Coding protocol and coder roles

51 A3. Data structure

52 A4. Extended evidence by theme

53 A5. Disconfirming and negative cases

53 A6. Limits of the Stage 1 evidence

54 **Appendix B: Stage 2: Longitudinal Textual Analysis**

54 B1. Panel Construction

56 B2. Dictionaries

57 B3. Relational-dictionary specification ladder

58 B4. Executional-dictionary specification ladder

59 B5. Shorter-window comparison

60 B6. External-dictionary and measurement validation

61 B7. Placebo, permutation, and survivor-bias checks

63 B8. Slope-break estimates by sample

63 B9. Robustness checks

64 B10. Ex-ante group construction and cross-measure heterogeneity

66 B11. Event-study and ex-ante group figures

67 B12. Directional changes in website language by organizational type

68 **Appendix C: Stage 3: Interviews and Roundtables**

68 C1. Design, sampling, and composition

68 C2. Interview and roundtable protocols

69 C3. Coding protocol and coder roles

69 C4. Evidentiary-status rules

70 C5. Data structure

70 C6. Outcome-independent configuration assignment and checks

72 C7. Extended evidence by configuration

76 C8. Mechanism evidence

77 C9. Disconfirming, reversed, and non-moving accounts

79 C10. Second-hand account register

79 C11. Sector coverage and evidentiary asymmetry

80 C12. Saturation, prevalence, and reflexivity

Appendix A — Stage 1: Executive Roundtable

A1. Design, facilitation, and participant composition

Stage 1 was a structured roundtable convened at a leading business school in June 2025 with 27 senior practitioners. Table 1 of the manuscript reports the composition. The session ran in four segments: a plenary introduction and round of self-introductions; a research-team presentation of the existing task-level evidence on GenAI and of a provisional grouping of organizations; two concurrent breakout groups; and a reconvened plenary in which each breakout reported back and the research team set out its intended next steps. Both rooms were recorded and transcribed verbatim. The two breakouts addressed, respectively, how different organizational types approach GenAI adoption, and how organizational structure, culture, and incentives mediate it.

The session was deliberately open-ended. Participants were told the research purpose at the outset and were asked to speak about how they expected GenAI to change knowledge work in their own sectors and about what they had already observed. The facilitators pressed on the difference between the two throughout, asking participants who described an anticipated shift what had happened so far, and asking those who described a change to state what had been done and when. Participants were explicitly invited to contest the working hypotheses and provisional groupings developed during the session, and did so. The provisional grouping concerned organizational types. The relational and executorial vocabulary, and the activity-level distinctions later developed in Stage 3, were not supplied to participants.

Chatham House rules applied. Participants consented to documenting the composition of the cohort, including disclosure of organizational affiliation in the composition tables, but not to attribution of specific remarks; quotations are therefore reported by role, sector, or session descriptor only.

Facilitation was shared between two members of the research team and two colleagues from the partner institute. This matters for the evidence base: contributions by members of the research team are treated as framing rather than data and are excluded from the evidence tables.

A2. Coding protocol and coder roles

Two of the authors met before coding began to agree the unit of analysis and the inclusion rules for marking passages. One coder was a senior strategy scholar with expertise in industry architecture and trained in qualitative analysis; the other was a researcher trained in qualitative analysis. They then coded the transcripts independently, working from the verbatim transcript rather than from notes, and compared their coding. The procedure followed Gioia, Corley and Hamilton (2013) as adapted for a single-session dataset:

- Pass 1 (open, in vivo). Each coder marked every passage in which a participant made a claim about what GenAI had changed, could change, or had failed to change, and labelled it in the participant's own language.

- Pass 2 (reconciliation). The two coders compared all marked passages, discussed singly marked passages, and agreed a common set of retained first-order codes. Differences concerned primarily where a passage began and ended; they were resolved by returning to the transcript.
- Pass 3 (second-order themes). Codes were grouped into candidate themes by each coder separately and then reconciled. Themes were required to be supported by at least two participants speaking independently; three candidate themes failed this test and were dropped, and are recorded in A5.
- Pass 4 (higher-order categories). Themes were grouped into the three higher-order categories reported in the manuscript. This step was performed jointly and was explicitly abductive: the categories were chosen for their bearing on the division of labor, which is the paper's question, not for their frequency in the transcript.

Because passage boundaries were not fixed in advance, a conventional inter-rater kappa would conflate segmentation and interpretation. We therefore report the nature and resolution of disagreements, the themes that did not meet the independent-support rule (A5), and the full first-order evidence (A4 and Table 2 of the manuscript).

A3. Data structure

Table A1. Stage 1 data structure

First-order codes (participant language)	Second-order theme	Higher-order category
Decomposition into tasks fails on "the work in between tasks"; project-leader role resists agent substitution; adoption concentrated in the layer the task account ranks as most exposed	The residual is the joining of tasks, not a harder task	The unit at which change is located
Generation cheap and fast; "good enough" content produced at volume; verification load displaced upward; gatekeeper capacity becomes the constraint	Cheap generation displaces effort onto verification	The unit at which change is located
Proposal to recast categories as scaling vs specialization tasks; question of whether the two can be separated within a role	Participants recast the categories at activity level	The unit at which change is located
Uptake immediate where the tool sits inside an existing platform; gradual for general-purpose tools; training targeted at the activity not the tool	Uptake follows the workflow, not the capability	The unit at which change is located
Audit demonstrably automatable yet unmoved; regulated professions expected to under-adopt; displacement described as inverse to licensing and control	Institutional accountability decouples capability from movement	Unevenness between capability and movement
Same activity adopted by corporate buyers and resisted inside provider firms; budget structure blocks a horizontal capability; funder-imposed cost ratios leave no budget line	Buyer structure, not the work, determines the outcome	Unevenness between capability and movement
Large firm negotiates data terms; SME boards cannot absorb the governance question; banned tools reappear as shadow use; regulated bank slows a rollout its staff demand	Scale determines the terms on which the exchange can be governed	Unevenness between capability and movement
Security, auditability, explainability described as the majority of deployment effort; realizing the capability requires organizational change in the buyer	Effort reallocates toward the conditions of reliance	Unevenness between capability and movement
Commoditized operation adopts fastest and proceeds are redirected upmarket; clients refuse to pay former prices; leverage pyramid described as shrinking	Repositioning within the existing bundle	The vocabulary of repositioning
Capital raised goes to distribution not compute; software differentiation attributed to the sales machine; model providers moving to industry-specific products	Technology-side repositioning around distribution and specificity	The vocabulary of repositioning
Relational work as advisory, trust, liability-bearing, legitimacy; executional work as analysis, process delivery, document production; counter-examples in both directions	Relational and executional as claims about defensible value	The vocabulary of repositioning

Note: n = 27 participants. First-order codes are compressed summaries of multiple in vivo passages; the underlying quotations are in Table A2 and in Table 2 of the manuscript. Only participant contributions are used as evidence.

A4. Extended evidence by theme

Table 2 of the manuscript reports the representative quotations. The table below reports further evidence not carried in the manuscript, retained here because several of these passages are the only support for a first-order code.

Table A2. Stage 1 supplementary evidence

Theme	Evidence	Role and sector
Adoption is bottom-up and driven by internal competition, not mandate	"We are not forced to use GenAI ... we are encouraged to use [it], but there is nothing if you don't use it. The reason there is very big adoption is because of the way the company is structured ... in order to live within this up-or-out world, you need to be better ... than the guy around the corner."	Partner, management consulting
Measurement of the gain is not achieved even where usage is measured	"We do observe the increase of the intensity of work. But that could be [for] multiple reasons ... I would say it is too early to attribute it to the tools."	Partner, management consulting
Generated output re-enters the firm's own knowledge base	"If that goes back into the archives, then we create this loop of training models on the stuff that models [have] created ... if the knowledge base starts getting polluted with generated content, maybe not so good."	Partner, management consulting
Professional identity bounds adoption	"Where I saw there was a difficulty in adopting GenAI is where you have a part of the organization that is around a mastery mindset, because people tended to feel that AI in general was like cheating."	Director, human resources strategy
Risk profile, not regulation, paces the rollout	"It is not even about the regulation, it is just a risk profile of the company ... sometimes we need to be slower, even if the appetite of our employees is give it to us now."	General manager, people analytics, banking
Governance capacity is the SME constraint	"The boards of those SMEs are not able to absorb into their governance framework how to govern this thing before they will unleash it to their organization."	Chair, science and innovation advisory
Where terms cannot be obtained, use moves out of view	"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."	Director, boutique advisory
Non-profits argued to be a distinct type	"I wouldn't categorize them as SMEs because they don't have a revenue ... any IP rights are not our rights ... I could not see any of that in those species that were mentioned."	Director of digital transformation, global health organization
Firm age proposed as the discriminating cut	"The long-term incumbents are, I think, probably more insulated than we imagine ... The ones that are challenged are the 20-year-old, 30-year-old companies that haven't actually had to do a pivot."	General partner, venture capital

A5. Disconfirming and negative cases

The following candidate themes were not carried forward for want of independent support, and are recorded here rather than suppressed: (i) that proprietary data would prove a durable differentiator, (ii) that GenAI adoption correlates with firm size, which participants split on in both directions; and (iii) that a distinct "wisdom work" tier sits above knowledge work, proposed by one participant and not taken up by others.

Two substantive disconfirming accounts are carried into the manuscript. The first is the claim that client expectations absorb productivity gains, so that leverage ratios move in the opposite direction to the one the pyramid-shrinking account predicts. The second is the observation that the layer most exposed on task-based measures is also the layer adopting fastest, which cuts against a simple displacement reading. Both are reported in Table 2 of the manuscript.

A6. Limits of the Stage 1 evidence

Stage 1 is a single session with a purposive sample assembled for variation across sectors and seniority. It supports claims about the categories and vocabulary participants used and about the range of positions in the room. It does not support prevalence claims, and no counts of participants endorsing a position are reported as evidence of prevalence.

Appendix B — Stage 2: Longitudinal Textual Analysis

Appendix B reports additional details for the Stage 2 website-language analysis: panel coverage, dictionary construction, fixed-effects specifications, external-dictionary validation, placebo and permutation tests, survivor-bias checks, robustness checks, ex-ante group construction, and descriptive website-language classifications.

B1. Panel construction

Sample construction began from the 626-firm Stage 2 sample, with each firm assigned to one organizational type and sample subsector before estimating language outcomes. For each firm-year, the scraper queried the Wayback Machine for the snapshot closest to July 1 of the target year, followed archived redirects where needed, retained the homepage and up to ten Guzman-and-Li keyword subpages, removed boilerplate and very short page captures, and aggregated retained text to the firm-year. The extended panel contains all usable 2017-2025 firm-years; the balanced pre-period specifications retain the 453 firms observed in each of 2017, 2018, and 2019; and the expanded unbalanced checks retain firms observed in at least one pre-period and one post-period year.

Table B1. Panel audit

Sample	Obs.	Firms	Years	PSF	ET	IT	TC
Core 2020-2025 panel	3,756	626	2020-2025	168	115	126	217
Extended 2017-2025 panel	5,330	626	2017-2025	168	115	126	217
Expanded 2020-2025 coverage before eligibility screen	8,289	2,122	2020-2025	479	973	127	543
Pre-period coverage	1,574	588	2017-2019	161	103	119	205
Balanced pre-period sample	4,077	453	2017-2025	126	72	100	155
Expanded survivor-bias estimation sample	7,221	1,502	2020-2025	350	613	127	412

Notes. PSF = professional services firms; ET = emerging technology; IT = incumbent technology; TC = traditional corporations. The balanced pre-period sample requires observations in each of 2017, 2018, and 2019. The survivor-bias estimation sample retains firms with at least one observation in 2020-2022 and at least one in 2023-2025; the larger expanded-coverage row precedes this restriction.

Table B2. Year-by-type coverage

Year	Organizational type	Firms	Median words	Mean rel.	Mean exec.
2017	Emerging Tech	76	5,315	178.8	103.7
2018	Emerging Tech	89	5,388	172.3	100.7
2019	Emerging Tech	98	5,734	165.9	115.4
2020	Emerging Tech	115	4,947	174.9	119.5
2021	Emerging Tech	115	5,095	177.6	130.0
2022	Emerging Tech	115	5,635	183.8	130.4
2023	Emerging Tech	115	5,969	187.5	141.8
2024	Emerging Tech	115	6,323	182.5	137.3
2025	Emerging Tech	115	6,793	193.4	157.8
2017	Incumbent Tech	105	6,030	191.5	105.1
2018	Incumbent Tech	110	6,618	186.5	104.1
2019	Incumbent Tech	118	6,584	189.3	114.9
2020	Incumbent Tech	126	5,522	182.1	112.6
2021	Incumbent Tech	126	6,050	197.6	120.1
2022	Incumbent Tech	126	6,274	189.7	120.5
2023	Incumbent Tech	126	6,360	189.3	128.8
2024	Incumbent Tech	126	6,546	194.3	142.3
2025	Incumbent Tech	126	7,002	199.2	149.6
2017	Professional Services	133	5,696	270.4	51.8
2018	Professional Services	142	6,136	268.0	55.4
2019	Professional Services	156	6,171	269.8	53.6
2020	Professional Services	168	5,770	280.0	54.7
2021	Professional Services	168	6,155	287.2	56.2
2022	Professional Services	168	6,328	301.1	55.4
2023	Professional Services	168	6,327	307.5	54.6
2024	Professional Services	168	6,716	322.9	57.8
2025	Professional Services	168	6,840	332.4	62.8
2017	Traditional Corporation	162	5,942	163.1	43.2
2018	Traditional Corporation	182	5,950	162.2	44.7
2019	Traditional Corporation	203	6,138	163.0	44.5
2020	Traditional Corporation	217	5,311	177.9	48.4
2021	Traditional Corporation	217	5,687	182.1	49.6
2022	Traditional Corporation	217	5,653	191.2	53.4
2023	Traditional Corporation	217	5,804	191.6	55.0

Year	Organizational type	Firms	Median words	Mean rel.	Mean exec.
2024	Traditional Corporation	217	6,012	200.9	54.2
2025	Traditional Corporation	217	5,699	203.7	58.0

Notes. Means are raw dictionary densities per 10,000 words.

B2. Dictionaries

The main Stage 2 table uses six dictionary constructions. The relational side includes a comprehensive relational dictionary and two constituent measures: client/advisory relational language and stakeholder-responsibility language. The executional side includes the comprehensive executional dictionary and two activity-oriented alternatives: interface/scalable-systems language and delivery/implementation language. All terms are matched using word boundaries on lowercased text and normalized per 10,000 words.

The same term lists are applied across the reported specification ladders. The comprehensive relational dictionary contains 50 terms, partitioned into 38 client/advisory terms and 12 stakeholder-responsibility terms. The comprehensive executional dictionary contains 50 terms; interface/scalable-systems and delivery/implementation language are measured with separate 29-term and 22-term dictionaries. These alternative measures assess the sensitivity of the estimates to the breadth and content of the vocabulary measured.

comprehensive relational: relationship, relationships, client, clients, partner, partners, partnership, partnerships, collaborate, collaboration, collaborative, trust, trusted, engage, engagement, engaged, support, supporting, together, commitment, committed, dedicated, advisor, advisors, advisory, tailored, guide, understand, team, insights, expertise, needs, goals, meet, network, strategic, leaders, growth, community, communities, governance, sustainability, sustainable, purpose, empower, empowering, transparency, inclusive, responsible, values

client/advisory relational: relationship, relationships, client, clients, partner, partners, partnership, partnerships, collaborate, collaboration, collaborative, trust, trusted, engage, engagement, engaged, support, supporting, together, commitment, committed, dedicated, advisor, advisors, advisory, tailored, guide, understand, team, insights, expertise, needs, goals, meet, network, strategic, leaders, growth

stakeholder responsibility: community, communities, governance, sustainability, sustainable, purpose, empower, empowering, transparency, inclusive, responsible, values

comprehensive executional: implement, implementation, implementing, execute, execution, executing, deliver, delivery, delivering, deploy, deployment, deploying, automate, automated, automating, automation, streamline, streamlined, streamlining, efficient, efficiency, efficiently, operate, operation, operations, operational, integrate, integration, integrating, scale, scalable, scaling, scalability, workflow, workflows, optimize, optimization, optimizing, optimal, accelerate,

accelerating, acceleration, orchestrate, orchestration, coordinate, coordination, coordinating, productivity, throughput, capacity

interface/scalable systems: automate, automated, automating, automation, workflow, workflows, orchestrate, orchestration, integrate, integration, integrating, platform, platforms, scale, scalable, scaling, scalability, standardize, standardized, standardization, streamline, streamlined, streamlining, optimize, optimization, optimizing, subscription, self serve, self service

delivery/implementation: implement, implementation, implementing, execute, execution, executing, deliver, delivery, delivering, deploy, deployment, deploying, operate, operation, operations, operational, coordinate, coordination, coordinating, capacity, throughput, productivity

Table B13 also reports two supplementary measures. Relationship language uses a narrower 24-term dictionary and is distinct from the 38-term client/advisory dictionary used in Table 3. Warrant/judgment language uses a separate 30-term dictionary. Both are normalized as occurrences per 10,000 words.

Relationship: relationship, relationships, client, clients, partner, partners, partnership, partnerships, collaborate, collaboration, collaborative, trust, trusted, engage, engagement, engaged, commitment, committed, dedicated, advisor, advisors, advisory, tailored, understand.

Warrant/judgment: accountability, accountable, assurance, assure, assured, counsel, discretion, discretionary, expert, expertise, guidance, independent, independence, integrity, interpret, interpretation, judgment, judgement, liability, liable, responsibility, responsible, trust, trusted, validate, validation, verification, verify, warrant, warranty.

B3. Relational-dictionary specification ladder

Table B3 reports the specification ladder for the three relational dictionaries used in the main Stage 2 table. Columns (1)-(4) report post-period level interactions under alternative fixed effects, type-specific trends, and pre-period coverage restrictions; columns (5)-(6) instead report slope interactions, with column (6) adding explicit GenAI-mention density.

Table B3. PSF relational estimates across dictionary constructions and specifications

	(1)	(2)	(3)	(4)	(5)	(6)
	Type + year FE	Firm + year FE	+ Type trends	Balanced + trends	Balanced slope break	+ GenAI control
Comprehensive relational	16.165 (0.029)	13.543 (0.080)	15.699 (0.085)	25.825 (0.007)	13.868 (0.005)	13.477 (0.006)
Client/advisory relational	12.121 (0.048)	9.624 (0.129)	7.753 (0.324)	15.059 (0.061)	9.138 (0.027)	8.809 (0.033)

	(1)	(2)	(3)	(4)	(5)	(6)
Stakeholder responsibility	4.044 (0.214)	3.918 (0.254)	7.946 (0.022)	10.766 (0.009)	4.730 (0.028)	4.668 (0.030)
Firm FE	No	Yes	Yes	Yes	Yes	Yes
Type FE	Yes	No	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Type-specific linear trends	No	No	Yes	Yes	Yes	Yes
Balanced pre-period sample	No	No	No	Yes	Yes	Yes
GenAI-mention density	No	No	No	No	No	Yes
Observations	5,330	5,330	5,330	4,077	4,077	4,077
Firms	626	626	626	453	453	453

Notes. Coefficients are reported with p-values in parentheses; values below .001 are shown as <.001. Standard errors are clustered by firm. Columns (1)-(4) report Professional Services x Post-2022. Columns (5)-(6) report Professional Services x Post-2022 slope. All models use the 2017-2025 panel. The balanced pre-period sample retains firms observed in each of 2017, 2018, and 2019. The reference organization type is traditional corporations. Columns (4)-(6) use the same balanced pre-period sample; the three coefficients in column (5) correspond to the three relational columns of manuscript Table 3.

B4. Executional-dictionary specification ladder

Table B4 reports the corresponding specification ladder for the three executional dictionaries used in the main Stage 2 table. The rows report emerging-technology and incumbent-technology coefficients, because these are the organizational types for which executional positioning is theoretically most relevant.

Table B4. Technology executional estimates across dictionary constructions and specifications

	(1)	(2)	(3)	(4)	(5)	(6)
	Type + year FE	Firm + year FE	+ Type trends	Balanced + trends	Balanced slope break	+ GenAI control
Comprehensive executional x ET	19.570 (<0.001)	19.701 (<0.001)	-0.309 (0.957)	-7.039 (0.313)	-1.865 (0.614)	-4.574 (0.217)
Comprehensive executional x IT	18.648 (<0.001)	19.701 (<0.001)	6.722 (0.282)	2.375 (0.733)	5.364 (0.130)	3.311 (0.364)
Interface/scalable systems x ET	19.552 (<0.001)	20.196 (<0.001)	2.362 (0.590)	0.187 (0.972)	1.605 (0.577)	-0.432 (0.883)
Interface/scalable systems x IT	17.512 (<0.001)	18.278 (<0.001)	5.302 (0.256)	3.197 (0.543)	2.722 (0.307)	1.178 (0.677)
Delivery/implementation x ET	1.394 (0.532)	1.032 (0.673)	-2.639 (0.429)	-5.507 (0.161)	-2.024 (0.297)	-2.738 (0.155)
Delivery/implementation x IT	0.406 (0.862)	1.083 (0.660)	-0.196 (0.949)	-2.820 (0.418)	2.030 (0.253)	1.489 (0.396)

	(1)	(2)	(3)	(4)	(5)	(6)
Firm FE	No	Yes	Yes	Yes	Yes	Yes
Type FE	Yes	No	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Type-specific linear trends	No	No	Yes	Yes	Yes	Yes
Balanced pre-period sample	No	No	No	Yes	Yes	Yes
GenAI-mention density	No	No	No	No	No	Yes
Observations	5,330	5,330	5,330	4,077	4,077	4,077
Firms	626	626	626	453	453	453

Notes. Coefficients are reported with p-values in parentheses; values below .001 are shown as <.001. Standard errors are clustered by firm. Columns (1)-(4) report Type x Post-2022. Columns (5)-(6) report Type x Post-2022 slope. All models use the 2017-2025 panel. ET = emerging technology; IT = incumbent technology. The omitted organization type is traditional corporations. Column (5) corresponds to the emerging- and incumbent-technology coefficients in the three executional columns of manuscript Table 3.

B5. Shorter-window comparison

Tables B5 and B6 report the corresponding fixed-effects ladders for the comprehensive relational and executional dictionaries. The 2020-2025 column is included as a shorter-window comparison; the main Stage 2 analysis uses the 2017-2025 panel.

Table B5. OLS predictors of comprehensive relational website language: fixed-effects ladder

	(1)	(2)	(3)	(4)	(5)
	2020-2025	2017-2025	2017-2025	2017-2025	Balanced
Professional Services x Post-2022	16.562 (0.024)	16.165 (0.029)	13.543 (0.080)	15.699 (0.085)	25.825 (0.007)
Emerging Tech x Post-2022	-5.960 (0.352)	-12.194 (0.066)	-11.372 (0.100)	6.625 (0.390)	9.599 (0.310)
Incumbent Tech x Post-2022	-10.532 (0.117)	-19.909 (0.004)	-19.831 (0.006)	2.907 (0.716)	0.675 (0.940)
Type FE	No	Yes	No	No	No
Firm FE	Yes	No	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Type-specific linear trends	No	No	No	Yes	Yes
Balanced pre-period sample	No	No	No	No	Yes
R ²	0.804	0.229	0.733	0.734	0.718
Observations	3,756	5,330	5,330	5,330	4,077
Firms	626	626	626	626	453

Notes. Coefficients are reported with p-values in parentheses; values below .001 are shown as <.001. Standard errors are clustered by firm. The omitted organization type is traditional corporations. Column (1) reproduces the core 2020-2025 post-period specification. Columns (2)-(4) use the extended 2017-2025 panel and vary the fixed-effects structure

and type-specific linear trends. Column (5) retains firms observed in each of 2017, 2018, and 2019. Dependent variables are dictionary terms per 10,000 words.

Table B6. OLS predictors of comprehensive executorial website language: fixed-effects ladder

	(1)	(2)	(3)	(4)	(5)
	2020-2025	2017-2025	2017-2025	2017-2025	Balanced
Professional Services x Post-2022	-2.339 (0.454)	-4.408 (0.144)	-4.259 (0.195)	-0.522 (0.873)	-2.497 (0.512)
Emerging Tech x Post-2022	13.761 (0.006)	19.570 (<0.001)	19.701 (<0.001)	-0.309 (0.957)	-7.039 (0.313)
Incumbent Tech x Post-2022	17.185 (0.003)	18.648 (<0.001)	19.701 (<0.001)	6.722 (0.282)	2.375 (0.733)
Type FE	No	Yes	No	No	No
Firm FE	Yes	No	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Type-specific linear trends	No	No	No	Yes	Yes
Balanced pre-period sample	No	No	No	No	Yes
R ²	0.850	0.251	0.810	0.812	0.812
Observations	3,756	5,330	5,330	5,330	4,077
Firms	626	626	626	626	453

Notes. Coefficients are reported with p-values in parentheses; values below .001 are shown as <0.001 . Standard errors are clustered by firm. The omitted organization type is traditional corporations. Column (1) reproduces the core 2020-2025 post-period specification. Columns (2)-(4) use the extended 2017-2025 panel and vary the fixed-effects structure and type-specific linear trends. Column (5) retains firms observed in each of 2017, 2018, and 2019. Dependent variables are dictionary terms per 10,000 words.

B6. External-dictionary and measurement validation

To assess whether the main dictionaries capture the intended language dimensions rather than a single general communication style, we compare them with Harvard IV-4 General Inquirer proxies and examine term-level structure within the scored panel. The Harvard IV-4 comparison uses independently defined word categories with limited overlap with the custom dictionaries. The measurement checks examine whether relational and executorial terms cohere more strongly within each measure than across the two measures. The Harvard IV-4 comparisons use the 2020-2025 panel of 626 firms and 3,756 observations. The relational proxy combines Affil and PowCoop (539 unique words); the executorial proxy combines Complet, Means, and Work, excluding words shared with the relational proxy (480 words).

Table B7. Harvard IV-4 General Inquirer validation

Validation	Statistic	Estimate
Harvard IV-4 relational proxy vs. custom relational	Correlation	0.531 (<0.001)
Harvard IV-4 executorial proxy vs. custom executorial	Correlation	0.463 (<0.001)
Harvard IV-4 relational proxy vs. custom executorial	Cross-correlation	-0.241 (<0.001)
Harvard IV-4 executorial proxy vs. custom relational	Cross-correlation	0.018 (0.262)
Post x PSF, Harvard IV-4 relational proxy	Panel regression	29.227 (0.034)
Post x ET, Harvard IV-4 executorial proxy	Panel regression	24.249 (0.069)
Post x IT, Harvard IV-4 executorial proxy	Panel regression	21.278 (0.111)

Notes. The analysis uses 626 firms and 3,756 firm-year observations from 2020-2025. Correlations are calculated at the firm-year level. Panel rows report organization-type-by-post-2022 level interactions from models with firm and year fixed effects, with traditional corporations as the reference type. These models do not include type-specific trends or post-2022 slope interactions. Regression p-values use two-sided normal-reference inference with firm-clustered standard errors and the finite-sample covariance correction used for the main models. P-values are in parentheses; values below .001 are shown as <.001.

Table B8. Measurement and discriminant diagnostics

Diagnostic	Statistic	Value
Within relational terms	Mean inter-term correlation	0.044
Within executorial terms	Mean inter-term correlation	0.091
Between relational and executorial terms	Mean inter-term correlation	0.006
Aggregate relational vs. executorial density	Correlation	0.043
Terms loading more strongly on PC1	PCA diagnostic	Relational: 12/50; executorial: 39/50

Notes. Diagnostics use 5,330 firm-year observations from 626 firms over 2017-2025 and the two comprehensive 50-term dictionaries. The PCA diagnostic reports how many of the 50 relational and 50 executorial terms load more strongly on the first principal component than on the second. Low between-measure correlations support treating the two dictionaries as distinct dimensions of positioning.

B7. Placebo, permutation, and survivor-bias checks

Table B9 reports supplementary checks of the 2020-2025 post-period models. First, the placebo-dictionary estimates are statistically imprecise for all three organizational-type interactions. Second, permutation tests shuffle organizational labels across firms to assess whether the focal coefficients are unusually large relative to randomly assigned sectors. Third, survivor-bias checks estimate the core 2020-2025 models on an expanded unbalanced panel of firms observed in at least one pre-period and one post-period year.

The placebo dictionary contains 50 generic corporate-communication terms, including company, annual, report, quarterly, board, director, employees, headquarters, global, press, conference, and awards; no placebo term overlaps with the relational or executional dictionaries. The permutation tests shuffle organizational-type labels across firms 1,000 times and re-estimate the same panel comparison on each draw; the reported p-values are randomization p-values from that null distribution.

Table B9. Placebo, permutation, and survivor-bias checks

Check	Statistic	Estimate
Post x PSF, placebo dictionary	Panel regression	-0.938 (0.928)
Post x ET, placebo dictionary	Panel regression	-3.086 (0.755)
Post x IT, placebo dictionary	Panel regression	-4.646 (0.670)
PSF relational, shuffled organizational labels	Permutation p-value	0.006
IT executional, shuffled organizational labels	Permutation p-value	<0.001
Professional services	Relational language	12.388 (0.028)
Emerging technology	Executional language	15.015 (<0.001)
Incumbent technology	Executional language	18.566 (<0.001)

Notes. All checks concern the 2020-2025 organization-type-by-post-2022 level-interaction models, with firm and year fixed effects and traditional corporations as the reference type; they do not directly test the post-2022 slope-break models. Placebo and permutation checks use 626 firms and 3,756 observations. Survivor-bias regressions use 1,502 firms and 7,221 observations, requiring at least one pre-period and one post-period observation per firm. Regression coefficients are followed by two-sided normal-reference p-values using firm-clustered standard errors and the finite-sample covariance correction used for the main models. Permutation values are the proportions of 1,000 label-shuffled coefficients at least as extreme in the observed coefficient direction. Values below .001 are shown as <.001.

B8. Slope-break estimates by sample

Table B10. Post-2022 slope-break estimates by sample

Sample	DV	Type	Coefficient	Obs.	Firms
Full 2017-2025	Relational	PSF	9.006 (0.071)	5,330	626
Full 2017-2025	Relational	ET	2.639 (0.568)	5,330	626
Full 2017-2025	Relational	IT	3.471 (0.396)	5,330	626
Full 2017-2025	Executional	PSF	2.051 (0.237)	5,330	626
Full 2017-2025	Executional	ET	1.006 (0.767)	5,330	626
Full 2017-2025	Executional	IT	5.706 (0.081)	5,330	626
Balanced pre-period	Relational	PSF	13.868 (0.005)	4,077	453
Balanced pre-period	Relational	ET	4.483 (0.408)	4,077	453
Balanced pre-period	Relational	IT	1.765 (0.705)	4,077	453
Balanced pre-period	Executional	PSF	1.204 (0.542)	4,077	453
Balanced pre-period	Executional	ET	-1.865 (0.614)	4,077	453
Balanced pre-period	Executional	IT	5.364 (0.130)	4,077	453

Notes. Coefficients are post-2022 slope changes from models with firm FE, year FE, and type-specific linear trends. P-values are in parentheses; values below .001 are shown as <.001. Standard errors are clustered by firm. Relational and executional refer to the comprehensive dictionaries. The reference organization type is traditional corporations. The balanced-sample estimates reproduce the comprehensive-dictionary columns of manuscript Table 3; the full-panel estimates use all 626 firms.

B9. Robustness checks

Table B11. Robustness checks for focal cells

Sample	Specification	DV	Type	Coefficient	Obs.
Full 2017-2025	Baseline	Relational	PSF	9.006 (0.071)	5,330
Full 2017-2025	Baseline	Executional	ET	1.006 (0.767)	5,330
Full 2017-2025	Baseline	Executional	IT	5.706 (0.081)	5,330
Full 2017-2025	Log word count	Relational	PSF	7.854 (0.108)	5,330
Full 2017-2025	Log word count	Executional	ET	0.960 (0.776)	5,330
Full 2017-2025	Log word count	Executional	IT	5.690 (0.082)	5,330
Full 2017-2025	Exclude 2020	Relational	PSF	8.229 (0.101)	4,704
Full 2017-2025	Exclude 2020	Executional	ET	1.079 (0.748)	4,704

Sample	Specification	DV	Type	Coefficient	Obs.
Full 2017-2025	Exclude 2020	Executorial	IT	5.583 (0.089)	4,704
Full 2017-2025	Winsorized 5/95	Relational	PSF	8.350 (0.036)	5,330
Full 2017-2025	Winsorized 5/95	Executorial	ET	-0.584 (0.811)	5,330
Full 2017-2025	Winsorized 5/95	Executorial	IT	2.115 (0.412)	5,330
Balanced pre-period	Baseline	Relational	PSF	13.868 (0.005)	4,077
Balanced pre-period	Baseline	Executorial	ET	-1.865 (0.614)	4,077
Balanced pre-period	Baseline	Executorial	IT	5.364 (0.130)	4,077
Balanced pre-period	Log word count	Relational	PSF	13.157 (0.006)	4,077
Balanced pre-period	Log word count	Executorial	ET	-1.900 (0.607)	4,077
Balanced pre-period	Log word count	Executorial	IT	5.388 (0.127)	4,077
Balanced pre-period	Exclude 2020	Relational	PSF	12.679 (0.011)	3,624
Balanced pre-period	Exclude 2020	Executorial	ET	-1.800 (0.617)	3,624
Balanced pre-period	Exclude 2020	Executorial	IT	5.210 (0.139)	3,624
Balanced pre-period	Winsorized 5/95	Relational	PSF	10.169 (0.018)	4,077
Balanced pre-period	Winsorized 5/95	Executorial	ET	-0.753 (0.793)	4,077
Balanced pre-period	Winsorized 5/95	Executorial	IT	2.526 (0.381)	4,077
Full 2017-2025	GenAI-mention density	Relational	PSF	8.754 (0.079)	5,330
Full 2017-2025	GenAI-mention density	Executorial	ET	-0.468 (0.892)	5,330
Full 2017-2025	GenAI-mention density	Executorial	IT	4.818 (0.146)	5,330
Balanced pre-period	GenAI-mention density	Relational	PSF	13.477 (0.006)	4,077
Balanced pre-period	GenAI-mention density	Executorial	ET	-4.574 (0.217)	4,077
Balanced pre-period	GenAI-mention density	Executorial	IT	3.311 (0.364)	4,077

Notes. Coefficients are post-2022 slope changes. P-values are in parentheses; values below .001 are shown as <.001. Standard errors are clustered by firm. Robustness specifications add log word count, exclude 2020, winsorize the dependent variable at the 5th and 95th percentiles, and add explicit GenAI-mention density. GenAI-mention density counts references to generative AI, GenAI, ChatGPT, OpenAI, large language models, LLMs, foundation models, copilots, and GPT per 10,000 words. All models include firm and year fixed effects, organization-type-specific linear trends, and organization-type-specific post-2022 slope interactions, with traditional corporations as the reference type. Relational and executorial refer to the comprehensive dictionaries.

B10. Ex-ante group construction and cross-measure heterogeneity

Table B12 reports the mapping used for the ex-ante heterogeneity analysis in the main text. The mapping uses the sample's subsector labels and is independent of the website-language outcomes. Table B13 applies the subgroup-versus-rest specification used in manuscript Table

4 to six language measures. Each estimate comes from a separate regression on the full 2017-2025 panel, comparing the focal subgroup with all remaining firms.

Table B12. Ex-ante group construction for heterogeneity analyses

Group	Assignment criterion	Included sample subsectors	Firms
Bundled advisory PSFs	Client-facing advisory models in which diagnosis, relationship management, and delivery are sold as an integrated engagement.	Management consulting; Investment banking; Creative & digital agencies	63
Accredited PSFs	Professional services in which formal credentials, legal responsibility, or regulated sign-off are central to the exchange.	Law firm; Accounting	73
Decomposable service providers	Service settings in which work is more often organized as separable deliverables, staffing, projects, or outsourced processes.	HR & staffing; Architecture; Engineering & environmental; IT services; BPO	28
Interface-oriented technology	Technology firms whose offerings are primarily scalable software, platforms, or data interfaces.	AI & data; Platforms & internet; Software & SaaS	100
Hardware/fintech technology	Technology firms whose offerings are tied more closely to infrastructure, devices, payments, or regulated transaction systems.	Semiconductors & hardware; Fintech & payments	44
Security/cyber technology	Technology firms whose offerings center on security, assurance, monitoring, and trust in digital systems.	Security & cyber	27

Notes. Groups are assigned mechanically from the frozen firm sample subsector labels before estimating the language models. Ambiguous, broad, or non-fitting subsectors are retained in the estimation sample but not reported as focal heterogeneity groups. The six focal groups contain 335 firms; the remaining 291 firms are retained in the comparison pool. A focal group's comparator also includes firms in the other five focal groups.

Table B13. Ex-ante group slope breaks across activity-relevant language measures

Group	Firms	Relational	Relationship	Warrant/judgment	Executorial	Interface	Delivery
Bundled advisory PSFs	63	12.849 (0.043)	6.341 (0.105)	2.192 (0.098)	4.536 (0.057)	-1.042 (0.631)	3.218 (0.020)
Accredited PSFs	73	2.570 (0.750)	-2.272 (0.730)	1.444 (0.374)	-2.227 (0.227)	-1.351 (0.205)	-0.458 (0.676)
Decomposable service providers	28	0.061 (0.994)	-4.528 (0.131)	1.147 (0.506)	-3.012 (0.538)	-4.315 (0.088)	0.414 (0.880)
Interface-oriented technology	100	3.170 (0.472)	1.328 (0.629)	1.374 (0.330)	9.970 (0.012)	8.189 (0.009)	2.274 (0.149)
Hardware/fintech technology	44	-3.808 (0.439)	-1.585 (0.461)	-1.568 (0.330)	-1.877 (0.630)	-1.226 (0.664)	-0.067 (0.978)
Security/cyber technology	27	1.818 (0.854)	-1.952 (0.721)	-6.726 (0.074)	-11.124 (0.046)	-5.623 (0.251)	-5.622 (0.068)
Firm FE		Yes	Yes	Yes	Yes	Yes	Yes
Year FE		Yes	Yes	Yes	Yes	Yes	Yes
Observations		5,330	5,330	5,330	5,330	5,330	5,330
Sample firms		626	626	626	626	626	626

Notes. Each cell reports the focal subgroup's post-2022 slope-change coefficient from a separate regression on the full 2017-2025 panel (626 firms; 5,330 observations), comparing that subgroup with all remaining firms. Models include firm and year fixed effects, a focal-group linear trend, and a focal-group post-2022 slope interaction. Coefficients measure changes in relative trends. Outcomes are term counts per 10,000 words. Relational and executorial denote the comprehensive dictionaries; interface denotes interface/scalable systems; delivery denotes delivery/implementation. Relationship and warrant/judgment are the supplementary dictionaries listed in Section B2. P-values are in parentheses; values below .001 are shown as <.001. Standard errors are clustered by firm. These are separate subgroup comparisons, not tests of pairwise differences between subgroup coefficients.

B11. Event-study and ex-ante group figures




Figure B1. Event-study coefficients, 2019 reference year.

Notes. Point estimates use the full 2017-2025 panel (626 firms; 5,330 observations). Each series uses the comprehensive dictionary indicated in the legend. Coefficients are organization-type-by-year interactions relative to traditional corporations and the 2019 reference year, with firm and year fixed effects. Confidence intervals are not displayed.

Post-2022 slope breaks by ex-ante subgroup




Figure B2. Selected ex-ante group slope-break estimates.

Notes. Selected estimates from manuscript Table 4. Each subgroup is compared with all remaining firms in the full 2017 - 2025 panel (626 firms; 5,330 observations), controlling for firm and year fixed effects and a subgroup-specific linear trend. Whiskers show 95% confidence intervals based on firm-clustered standard errors. ** $p < .05$; *** $p < .01$.

B12. Directional changes in website language by organizational type

Table B14. Directional changes in website language by organizational type

Type	Both	Executional	Relational	Neither
Emerging Tech	53 (46.1%)	27 (23.5%)	18 (15.7%)	17 (14.8%)
Incumbent Tech	46 (36.5%)	35 (27.8%)	26 (20.6%)	19 (15.1%)
Professional Services	67 (39.9%)	27 (16.1%)	42 (25.0%)	32 (19.0%)
Traditional Corporation	87 (40.1%)	36 (16.6%)	48 (22.1%)	46 (21.2%)

Notes. Directional classifications compare average 2023-2025 language with average 2020-2022 language at the firm level. These are descriptive changes in public website language and do not measure internal reorganization or activity reallocation.

Appendix C — Stage 3: Interviews and Roundtables

C1. Design, sampling, and composition

Stage 3 comprises 33 semi-structured interviews conducted between late 2025 and early 2026 with senior practitioners, and two structured roundtables held in New York (Roundtable A) and San Francisco (Roundtable B) with 40 participants in total. Table 5 of the manuscript reports the composition. Interviews ran 30 to 75 minutes; roundtables ran approximately five hours each, in plenary and breakout formats.

Sampling was purposive and sought variation on three dimensions: organizational type (professional services, traditional corporation, incumbent technology, emerging technology); regulatory intensity of the activity discussed (from statutory audit to marketing copy); and position relative to the exchange (provider, buyer, investor, regulator-adjacent). Interviewees were drawn from the membership of a UK directors' institute, which biases the sample toward UK and Channel Islands respondents, C-level executives and/or corporate directors, and toward smaller and mid-sized organizations; the roundtables were US-based and biased toward large technology firms and investors. The two sources are therefore complementary in coverage and are reported separately throughout.

C2. Interview and roundtable protocols

Interviews followed a protocol of four blocks, asked in order but not read verbatim, so that informants described activities in their own terms before being asked about allocation and exchange:

- **Block 1 — Production.** What has GenAI changed in what your organization actually does day to day? Where has it worked and where has it not? Probe for the specific activity, not the tool.
- **Block 2 — Allocation.** For each activity named: who performs it now, who performed it before, and when did that change? Probe: has anything stopped being sold, bought, or offered?
- **Block 3 — Exchange.** How is that activity specified, priced, verified, and warranted? What would have to be true for it to be bought from someone else, or brought in house?
- **Block 4 — Competition and entry.** Who is gaining and losing? What kinds of organization are entering? What is the binding constraint on them?

Roundtables were organized around two framing questions - who does what, and who captures the value - with breakouts on services, software, manufacturing, and moats. Facilitators pressed participants for concrete referents whenever a general claim was made.

C3. Coding protocol and coder roles

Coding proceeded from the activity outward rather than from the informant or the theme. The unit of analysis is the activity-account: one informant's description of one piece of knowledge work. The same two coders from Stage 1 worked independently through the corpus in three passes.

- **Pass 1 — Extraction.** Every passage naming a specific piece of knowledge work was extracted with four fields recorded verbatim where possible: the activity; what was said to have changed about its production; who performs it now; and what the informant identified as producing the change. Passages describing tools or sentiment without a named activity were not extracted.
- **Pass 2 — Status assignment.** Each account was assigned an evidentiary status under the rules in C4, independently by both coders. Differences concerned what constituted a single activity rather than the evidentiary weight of an account; they were reconciled by returning to the transcript and applying the identifiable-referent test.
- **Pass 3 — Characteristic coding.** Each activity was coded on the two framework dimensions using task and exchange characteristics only, with the reported allocation outcome masked (C6).

The two coders brought complementary vantage points to the material. To keep coding grounded in the record, reconciliation was conducted against transcribed evidence, and the identifiable-referent test in C4 was applied to what appeared in the transcript rather than to background knowledge about the case.

C4. Evidentiary-status rules

Table C1. Evidentiary status definitions and decision rule

Status	Definition	Test applied
Realized	A change in who performs the activity, or in how it is bought, that has already occurred	The transcript contains an identifiable referent: a named service line no longer offered, an operations center restructured, a role removed, a price actually changed, a contract actually signed, or a budget actually not reallocated
Realized, reversed	A realized change that was subsequently undone	As above, plus an explicit statement that the activity returned
Realized (non-movement)	The informant states affirmatively that the activity has not reallocated, and identifies the constraint	The non-movement is asserted, not inferred from silence
Prospective	An expectation, intention, or decision not yet implemented	No identifiable referent in the past tense; future or conditional framing
Second-hand	A realized change reported about another organization	Speaker is not a member of the organization in which the change occurred
Non-moving / persistence	The activity is described as unchanged and expected to remain so	Explicit statement of persistence with a reason

Note: Claims about reallocation in the manuscript rest only on accounts coded Realized or Realized, reversed. Silence about an activity is never coded as non-movement.

C5. Data structure

Table C2. Stage 3 data structure

First-order codes	Second-order theme	Analytical role
Output generated without mobilizing the expertise formerly required; analyses arriving in hours that took teams weeks; expert panel judgment encoded into a product; delivery inverted to machine-primary	Production without the embodied expert	Framework dimension: expertise bypass
Framing named as the limit across sectors; interaction between obligations not handled; exceptions more costly per case after automation	Bypass stops at framing and at linking	Framework dimension: expertise bypass
Project re-specified as a subscription; annual engagement replaced by continuous provision; bundling toward a monthly managed service	Episodic engagement converted to a standing interface	Framework dimension: exchange governability
Contingent pricing where the outcome is countable; disputes over whether outcomes can be specified at all; validation purchased separately from production	Specifiability and verifiability determine the contract form	Framework dimension: exchange governability
Sign-off, indemnity insurance, accreditation, statutory requirement; a named party to hold responsible	Warrant cannot be assigned to a machine	Framework dimension: exchange governability
Shared design platforms merging independently produced components; open banking and ledger interoperability; data dictionaries and semantic layers; certification thresholds for suppliers	Complementary infrastructure, much of it predating GenAI	Complementary condition
Integration across heterogeneous systems declined by platform owners; connector-building ventures; scarcity of accredited AI auditors	Who supplies the conditions of exchange	Corollary: entrepreneurial position
Acquisition of the client relationship where the market resists; federated specialist teams; very small teams competing on analytical output	Entry forms opened by the configurations	Corollary: entrepreneurial position

C6. Outcome-independent configuration assignment and checks

Activities were assigned to configurations from characteristics coded independently of their reported allocation outcome. Both coders first coded the two framework dimensions from transcribed activity and exchange descriptions. To preserve outcome independence, the characteristic-coding sheet omitted the allocation and evidentiary-status fields; disagreements on the two dimensions were reconciled against transcribed evidence before the configuration was mechanically derived. Expertise bypass was coded high where the focal output could be generated without mobilizing the embodied expertise formerly required. Governability was coded high where the account established all four of: manageable interdependence, ex-ante specifiability, ex-post verifiability, and assignable warrant or liability. Where an account distinguished machine-generable production from accountable warrant, sign-off, or professional judgment, the focal warranting activity was coded separately because that was the activity whose allocation was at issue.

The table below reports every activity in Table 6 of the manuscript with its codes on both dimensions and the resulting configuration, so that readers can see how the coded characteristics combine into configurations.

Table C3. Configuration assignment checks

Activity	Expertise bypass	Exchange governability	Configuration
Expert identification in a therapy area	High	High	1
Real-world data source selection and evidence planning	High	High	1
Payer research	High	High	1
Security operations center monitoring	High	High	1
Penetration testing	High	High	1
Literature screening and extraction	High	High	1
Medical publications content development	High	High	1
Analytics and insight production	High	High	1
Graphic and visual asset production	High	High	1
Cross-firm invoicing and proof-of-delivery retrieval	High	High	1
Bill-of-materials procurement and value engineering	High	High	1
Marketing copy and content drafting	High	Resistant	2
Problem framing and question reframing	Low	Resistant	4
Market research for client engagements	High	Resistant	2
Multi-tax planning where obligations interact	Low	Resistant	4
Bookkeeping and year-end preparation	High	Resistant	2
Lower-tier analysis in government advisory	High	Resistant	2
Employee wellbeing diagnosis and action planning	High	Resistant	2
Consultancy delivery inverted to machine-primary	High	Resistant	2
Board packs and financial analysis	High	Resistant	2
Sales training and capability transfer	High	Resistant	2
Publication support tasks	High	Resistant	2
Exception handling in accounts payable	High	Resistant	2
Inter-firm design coordination	Low	High	3
Independent validation and gap analysis	Low	High	3
Technical and production architecture drawings	Low	High	3
Look-back tax recovery	High	High	1
Contracted delivery of IT services	Low	High	3
Temporary specialist and freelance provision	Low	High	3
Network-level scam interdiction	High	Resistant	2
Statutory audit opinion and accountable assurance	Low	Resistant	4
Audit of machine-produced work and of models	Low	Resistant	4
Legal advice requiring professional sign-off	Low	Resistant	4
Independent second-opinion verification	Low	Resistant	4
Accountability-bearing sign-off on translated content	Low	Resistant	4
Regulatory approval and building control assessment	Low	Resistant	4

Note: Configuration 1 = contestable reallocation; 2 = persistent bundles and selective insourcing; 3 = federated provision of human expertise; 4 = integrated expert provision. High bypass means the focal activity's output could be generated without mobilizing the embodied expertise formerly required; low bypass includes accountable warranting, sign-off, framing, and professional judgment activities even when adjacent research or drafting steps are machine-generable.

C7. Extended evidence by configuration

The manuscript reports one or two quotations per configuration. The table below reports the remaining evidence, so that each configuration can be assessed on more than its illustrative case.

Table C4. Supplementary activity-level evidence by configuration

Config.	Activity and account	Evidence
1	Expert identification ceased to be a purchasable engagement (pharma advisory)	"There's actually, from [a] strategic perspective, no point in us offering to a client, we'll do expert identification, because they say, well, why would we ask you to do that? We can just buy it." Qualifier from the same informant: "There may well be companies who are maybe not quite as big ... where for a number of years, you will still be able to win projects with them."
1	Evidence planning converted from project to subscription (pharma advisory)	"What we would do in the past is we would do that as a consultancy project ... from a blank sheet of paper. That has now increasingly become a platform-based project, a subscription project."
1	Payer research interrogated from an archive (pharma advisory)	"You can build a payer avatar based on all of that stored information in order that someone can actually ask it questions or test their messages through this AI-driven archive."
1	Look-back tax recovery under contingent pricing (Roundtable B)	"Let us have the engine go through every single line and transaction ... And the outcome is I get 40 cents on the dollar." Contested at the same table: "Outcome so far is a sales trick to me."
1	Penetration testing converted to continuous provision (cybersecurity)	"What AI is allowing us to do ... is continuous pen testing ... which an AI can do, but obviously a person can't." Qualifier: "You still get the people who do it once a year, because they need to for government authorization, and they'll do the bare minimum."
1	Bundling toward a standing exchange (cybersecurity)	"Instead of paying us once in a while to do certain pieces of work, we want to say, pay us 20k a month ... you'll actually see quite a change, not in the products itself, but how it's packaged."
2	Market research repriced and shifted toward client-side provision (management consulting)	"You were charged 50k a pop for something like that ... maybe that should now be five, or 10. And maybe the client should expect it in a week, not a month ... you can advise them to insource the capability ... they can employ fairly junior analysts."
2	Wellbeing action planning moved to the client's HR team (technology advisory)	"That was always the barrier ... they had to engage a third party to deliver some solutions. Now, the HR team, generally, are empowered to then take some actionable steps ... it's less important that it's 100 per cent accurate."
2	Bookkeeping re-integrated into a small practice (accounting)	"Whereas three years ago, we didn't really do bookkeeping. We are taking on a lot more bookkeeping clients because it's cost-effective for us to do it because of the ease with software." Attributed by the informant to accounting software and open banking, not GenAI.

Config.	Activity and account	Evidence
2	Capability transfer with an explicit hand-off horizon (sales consultancy)	"We are there to facilitate their learning so they can take it in-house ... maybe six quarters, eight quarters, to embed the culture and the learning ... and then they can hopefully run it themselves."
2	Outsourced publication support absorbed internally (governance)	"I ended up reducing headcount by one because ... I found tech tools that would do the work around it and tech tools that would help me write better. And I didn't actually need that role anymore."
2	Exceptions become more costly per case (financial transformation)	"You have automated 95 per cent of the process ... this working on this remaining 5 per cent, which is not falling to the pattern ... will take them probably much more time than it used to be."
3	Federated teams competing with integrated firms (construction)	"Instead of the big guy offering a one-stop shop package to a client ... I can offer almost the same performance and delivery ... because everything's being done using cloud-based tools, and it's all integrated into a single model." Prior state: "A lot of manual checking."
3	Validation purchased separately from production (real estate and infrastructure)	"They'll create the goods, and they'll give them to me to validate. I will then go back to them and say, I don't understand that. That's not correct ... And do a gap analysis."
3	Technical drawing separated on liability grounds (construction)	"There's a lot of creative architects now who don't want to do that piece because it carries liability with it ... they've set up a division of their company that just does these technical drawings ... That has nothing to do with AI. That has everything to do with contract law and who's responsible."
3	Freelance provision made cheaper to assemble (management consulting)	"A lesser permanent staff and work with freelancers ... because you have a knowledge base, your structured data, your in-house AI, companies can do it."
4	Audit demonstrably automatable and unmoved (financial services)	"The job of audit can be autonomously done. Can be done, right? ... But I have not seen any evidence of any audit firm, or any entrepreneurial company ... actively transforming ... the sector that I'm working [in], which is audit."
4	The binding constraint named as assurance, not capability (financial services)	"Our challenge with my startup is how to provide the evidence and the assurance to a particular customer that an AI doing it autonomously can be reliable without a human."
4	Regulated professionals expected to under-adopt (financial services governance)	"Probably the sectors where you've got professional requirements ... solicitors, accountants, auditors ... won't get sucked into this quite so much because their personal and professional reputations rely on it and their insurance relies on it."
4	Demand for second opinion rising (accounting and tax)	"They just want to check that what they've seen on the internet is correct."
4	Audit of models not presently performable (Roundtable A)	"There are very few people who are authorized to audit AI. So, there are only, apparently in the United States, 70 people." And: "The function of audit is largely human today. It's not being productized."
4	Warrant could not be assigned to a machine (Roundtable B)	"I needed somebody when things go wrong to call on and blame. I cannot blame that Google translation did this to me." Same table: "The throat to choke is solidly underestimated."

Table 6 of the manuscript groups the 36 activity-level accounts by configuration and evidentiary status. Table C5 reports each activity, sector, described change, evidentiary status, and qualifying or disconfirming account. Thirteen of the 36 accounts include qualifying or disconfirming evidence; three of these rows include explicit attribution to changes other than GenAI.

Table C5. Complete activity-level inventory (all columns)

Activity	Sector	Change described, and who now performs it	Evidentiary status	Qualifying or disconfirming accounts
Expert identification in a therapy area	Pharma advisory	Bespoke consultancy search replaced by subscription databases; no longer offered by the provider	Realized	Buyers later in the adoption curve continue to commission the project (same informant)
Real-world data source selection and evidence planning	Pharma advisory	Blank-sheet consultancy project replaced by subscription access to an aggregated, machine-interrogable archive	Realized	-
Payer research	Pharma advisory	Repeated primary research replaced by interrogation of a stored archive, including simulated payer respondents	Realized	-
Security operations center monitoring	Cybersecurity	Human analyst teams replaced by proprietary models under expert oversight; technical headcount reduced by about two-thirds through attrition	Realized	-
Penetration testing	Cybersecurity	Annual point-in-time engagement extended to continuous machine-run testing; sold as a new product line	Realized	Buyers with minimum compliance requirements continue to purchase annual tests (same informant)
Literature screening and extraction	Pharma advisory	Protocol-governed screening reallocated to in-house machine execution; specialist platforms emerging as external substitutes	Realized reallocation	-
Medical publications content development	Pharma advisory	Client moved the activity to a large technology platform on consultancy advice and sought a 30 per cent price reduction; the platform did not perform and the agency resumed the work	Realized, reversed	-
Analytics and insight production	Cross-sector (data provider)	Clients insourced production, then returned it to the provider within about six months	Realized, reversed	-
Graphic and visual asset production	Defense and security consulting	Work previously contracted to a designer now self-served	Realized	-
Cross-firm invoicing and proof-of-delivery retrieval	Construction distribution; accounting	Manual retrieval across carriers, and manual invoice entry, replaced by automated channels	Realized	Both informants attribute the change to automation and open banking rather than to GenAI
Bill-of-materials procurement and value engineering	Construction supply	Complete digital specification would permit agent-led procurement, bypassing the value-engineering intermediary	Prospective	-
Marketing copy and content drafting	Marketing consultancy	Two clients now produce first drafts; the provider corrects to house style and supplies strategy	Realized, partial	Provider expects the correction activity to move to the client as well (same informant)

Activity	Sector	Change described, and who now performs it	Evidentiary status	Qualifying or disconfirming accounts
Problem framing and question reframing	Strategy and management consulting	Described as not bypassed, and as the first step of every engagement	Realized (persistence)	-
Market research for client engagements	Management consulting	Market research shifted toward client-side provision using junior analysts; price described as falling from about GBP 50,000 to GBP 5,000-10,000 and delivery from a month to a week	Realized (client-side shift)	-
Multi-tax planning where obligations interact	Accounting and tax	Machine answers described as accurate on single taxes but not on their interaction; the focal planning activity remains with the provider.	Realized (persistence)	A defense consultant reports machine-generated UK tax guidance that he has never found an accountant to contradict
Bookkeeping and year-end preparation	Accounting	Automation lowered unit cost sufficiently that the accountancy practice absorbed bookkeeping previously performed elsewhere	Realized, re-integration	Attributed to accounting software and open banking rather than to GenAI
Lower-tier analysis in government advisory	Government advisory	Client-side lower-tier analysis described as machine-performed at levels one to three, sometimes four; human judgment enters at four to six	Realized (client-side shift)	-
Employee wellbeing diagnosis and action planning	Technology advisory	Expert panel judgment encoded into the product; the client's HR team now performs action planning previously bought from a third party	Realized	Accepted because the cost of an imperfect answer is low (same informant)
Consultancy delivery inverted to machine-primary	Technology and business advisory	A fifteen-consultant firm reported as having inverted its delivery model and replicated it across countries	Realized (second-hand)	Reported by a vendor about a client rather than observed directly; very small firm
Board packs and financial analysis	Cybersecurity	Attempted in-house; output not yet judged satisfactory	Prospective (attempted)	-
Sales training and capability transfer	Sales consultancy	Capability transfer to the client under way; provider building client capability toward intended insourcing within six to eight quarters	Realized (client-side shift)	-
Publication support tasks	Governance and company secretarial	Previously outsourced tasks absorbed internally; one role removed	Realized	-
Exception handling in accounts payable	Financial transformation	Accounts-payable work shifted toward client-side handling; routine automation leaves residual exceptions more time-consuming per case	Realized (client-side shift)	-
Inter-firm design coordination: clash detection, model merging, error resolution	Construction	Manual checking across independently produced components replaced by tools embedded in shared design software; federated teams now compete with integrated firms	Realized	-
Independent validation and gap analysis of consultant deliverables	Real estate and infrastructure	Client contracts an external subject-matter expert to validate machine-assisted output produced by a large consultancy's junior team	Realized	-
Technical and production architecture drawings	Construction	Separated into a specialist division and traded between firms because of how liability is assigned	Realized	Informant attributes the change to contract law and explicitly not to AI
Look-back tax recovery	Tax advisory	Machine review of every transaction enables contingent contracting on amounts recovered, at about forty cents on the dollar.	Realized reallocation	-

Activity	Sector	Change described, and who now performs it	Evidentiary status	Qualifying or disconfirming accounts
Contracted delivery of IT services	IT services	Shift described from time-and-materials to outcome-based pricing	Realized reallocation	Two participants at the same roundtable dispute that outcomes can be specified outside standardizable work
Temporary specialist and freelance provision	Management consulting	Structured knowledge bases and in-house models reduce the cost of onboarding temporary collaborators	Realized	Buyer certification and creditworthiness thresholds exclude small providers (2 accounts)
Network-level scam interdiction	Telecommunications	Capability demonstrated; the service could not be priced because the outcome could not be expressed in a contractible metric	Realized (non-contracting)	-
Statutory audit opinion and accountable assurance	Accounting	Autonomous execution of audit procedures demonstrated at prototype; no reallocation of the statutory audit opinion observed; the binding constraint named as the inability to evidence assurance to the buyer.	Realized (non-movement)	One entrant argues that incumbent assurance rests on sampling and is therefore contestable
Audit of machine-produced work and of models	Assurance	Cannot presently be performed: model observability described as poor to absent, and about seventy people in the United States authorized to audit AI systems	Realized (constraint)	-
Legal advice requiring professional sign-off	Law	Internal adoption described as fast, including precedent search and paralegal work; accountable legal advice and professional sign-off do not reallocate.	Realized	One informant reports law firms as the fastest adopters he has observed, contrary to his own expectation
Independent second-opinion verification	Accounting; strategy and defense consulting	Demand described as rising, because buyers arrive with machine-generated answers they cannot themselves warrant	Realized	-
Accountability-bearing sign-off on translated content	Large technology firm	A ten-million-dollar budget was not reallocated to machine translation over six to seven years, for want of an accountable party	Realized (non-movement)	-
Regulatory approval and building control assessment	Construction	Regulator has not adopted; assessment throughput constrains the rest of the chain	Realized (constraint)	-

Note: This table reproduces in full the inventory summarized in Table 6 of the manuscript. A dash indicates that no qualifying or disconfirming account was recorded for that activity.

C8. Mechanism evidence

Table 7 of the manuscript classifies mechanisms into three classes. The purpose of the classification is to prevent attributing all standardization and coordination change to GenAI. The table below supplies the evidence for the Class C mechanisms in particular, since those are the ones the manuscript relies on to disclaim attribution.

Table C6. Evidence for complementary-infrastructure mechanisms (Class C)

Mechanism	Evidence	Attribution stated by informant
Structured repositories and data ontologies	"You have to build a data dictionary or ontology ... defining each term that's being used on the system, and the banks are already investing very, very significant amounts of money in building these data dictionaries."	Precondition for AI, not produced by it
Rising data-preparation burden	"Two years ago ... about a month worth of ... cleaning up the data. And then over the last two years ... 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."	Prerequisite, and increasing
Shared cloud design and modelling platforms	"There's some AI tools built into the Autodesk Revit software to help them merge that stuff and deal with all the things like clash detection, errors detection."	Embedded in existing design software
Open banking and ledger interoperability	"If you don't have open banking, none of these software will work." And: "If I've got an invoice to one of my clients that's also on [the same package], I can email it straight into their bookkeeping software ... there's no human effort needed."	Explicitly not GenAI
Certification and procurement infrastructure	"If you haven't got 27001, you're going to struggle." And: "It gives us more critical mass ... they've got a balance sheet which is sufficiently robust for us to deal with them. So it's a creditworthiness aspect."	Independent of capability
Systems integration and authentication across heterogeneous systems	"Can you integrate across all the case management systems, across all the different agencies? ... The short answer from OpenAI is, oh, that's your job, go call [a systems integrator]." And: "Can you integrate your single sign-on ... across all the tools and the case management systems ... and [the vendor] is like, oh yeah, we can't do that."	Declined by platform owners
Scale as a condition of governable terms	"We have an agreement with [the model provider] that none of our data will be used to train any future models ... We can do that because we are a huge global organization with generally quite deep pockets. Would a mid-cap firm be able to have that negotiating power?"	Bargaining power, not technology

C9. Disconfirming, reversed, and non-moving accounts

To make contrary and qualifying evidence visible, the table below reports every account that cuts against a directional claim in the paper, together with what it disconfirms.

Table C7. Disconfirming and contrary accounts

Account	What it cuts against	Sector and role
"I've not seen anybody, in a large organization, make a case for restructuring and rationalizing the workforce based on the availability of LLMs."	Any claim that GenAI is driving observable workforce reorganization	Management consulting, partner
"There's no correlation between AI and employment that I can see from my chair ... [firms] are right-sizing and they're pointing to AI so they don't get into trouble from a perception point of view ... it's an excuse."	The inference from headcount reduction to AI substitution	Professional and technical services, advisor
"In terms of any breakthrough new products or services, no, not yet, nothing ... I'm not seeing game-changing disruption yet."	The claim that GenAI has produced new offerings	Diversified boards, director
On proprietary data as a differentiator: "I've not seen that, and I'm surprised." And, of a vendor's proprietary-data marketing: "I think that's marketing."	The proprietary-data thesis the research team's own earlier work reported	Diversified boards; accounting
"We have more junior people per senior people today than we did 15 years ago, because client expectations ... have gone through the roof faster than the technology has been able to keep up."	The shrinking-pyramid account	Leadership advisory, partner
Machine-generated UK tax guidance "which so far I've never found any accountant or tax person [say] that was wrong."	The multi-tax interaction limit claimed by an accountant	Defense and security consulting, managing director
"I expected [law] to be the slow part, actually. So I'm surprised that they are fast."	The expectation that regulation slows adoption	Private equity and law, managing director
Headcount rising, not falling, in procurement and supply chain: "they're not reducing the number of people ... they're having to bring in more people."	Displacement as the general case	Procurement and supply chain, independent verifier
"Generative AI has basically destroyed the ad industry ... what you're seeing is mergers between all the really big guys."	The claim that no sector shows realized structural change; runs the opposite way	Brand strategy, managing director
On differentiation: "I don't think it helps with differentiation. I think it creates vanilla."	The claim that GenAI raises the ceiling of output quality	Brand strategy, managing director
Medical publications content development returned to the agency after the platform "didn't work at all. It was useless."	The durability of Configuration 1 reallocation	Pharma advisory, COO
Analytics production insourced and returned within about six months: "actually, no, that's pretty hard."	The durability of client insourcing	Enterprise data and analytics, Roundtable A

C10. Second-hand account register

Three accounts used in the manuscript are second-hand. They are listed here in full because the evidentiary weight they can bear is limited.

Table C8. Second-hand accounts

Account	Reported by	Limitation
A fifteen-consultant advisory firm inverted its delivery model so that "80 per cent of the work is now done by AI", and replicated it across countries	A vendor, about a client firm	Not observed directly; very small firm; the vendor sells the capability described
Approximately 70 people in the United States authorized to audit AI systems	An investor, about a portfolio company	Hedged by the speaker as "apparently"; not independently verified
A pharmaceutical company was advised by a consultancy to move content development to a technology platform and to seek a 30 per cent price reduction	The displaced agency	Reported by the party that lost and then regained the work

C11. Sector coverage and evidentiary asymmetry

Coverage is uneven and the unevenness is not random. Professional services are heavily represented and traditional corporations are not; the activities on which we have realized accounts are concentrated in advisory, accounting, construction, and cybersecurity; and we have no realized account from a large regulated financial institution, which is the setting in which the governability constraint would be most severe. Claims about Configuration 4 therefore rest more on non-movement and stated constraint than on observed change, which is a weaker evidentiary base than Configurations 1 and 3.

One informant volunteered an alternative interpretation of the Stage 2 finding, and it is reported here rather than answered: asked whether the relational shift in website language was related to GenAI, she replied that "it could be that people ... [had] to think more with a sharper lens on what is it that they bring that is different and that is special", because "you can't actually charge them the ... big four premium for doing that kind of stuff", and added - unprompted - "maybe they're actually using AI to help them to find language that is going to differentiate. That would be funny, wouldn't it?" We treat this as the strongest available statement of the marketing-and-legitimacy interpretation of Stage 2, offered by a participant rather than constructed by us.

C12. Saturation, prevalence, and reflexivity

We do not claim theoretical saturation. The setting changed materially during the fieldwork window, the sample was assembled for variation rather than for estimation, and new activity types were still appearing in the final interviews. We do not convert account frequency into prevalence, and the counts reported in Table 6 of the manuscript are counts of accounts, not of firms or of incidence.

To keep the evidentiary roles separated, passages in which Stage 2 results were introduced before eliciting informant reaction were coded as commentary rather than as independent support for the Stage 2 interpretation. The exception noted in C11 is retained because it runs against the reported reading. C9 reports the contrary and qualifying evidence retained in the analysis.

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.