

Epistemic Governance as Infrastructural Practice: Reframing the Structure-Agency Relationship in AI-Augmented Knowledge Work

Dr Suela Pirushi

Authors' note. The author thanks the 127 professionals across six organisations who participated in this research. This study forms part of a broader doctoral research programme on epistemic governance in AI-augmented professional services. A methodological supplement (coding tree, power-quote table, theoretical saturation matrix, sampling decision log, and reflexivity notes), interview protocols, and supporting materials are available at <https://osf.io/fw5kj> (<https://doi.org/10.17605/OSF.IO/FW5KJ>). Correspondence concerning this article should be addressed to suelapirushi@yahoo.com.

Abstract

Organisational theory has long been divided between structural accounts of knowledge governance and agentic accounts of professional judgment. Artificial intelligence (AI) unsettles both: it severs the link between production and accountability while generating the very fluency markers professionals use to assess quality. Through comparative case analysis of six professional services organisations (127 interviews, 43 documents, and 22 observation sessions over 11 months), this study develops a grounded theory of epistemic governance through infrastructuring, the deliberate design of background systems that make the epistemic status of knowledge claims visible and actionable. The study describes a confidence-decoupling phenomenon, a self-reinforcing pattern in which AI-enhanced surface quality raises organisational confidence even as epistemic quality declines, connecting it to March's (1991) exploitation trap. It develops a three-path process model showing how leaders' framing of epistemic shocks is associated with whether organisations build infrastructure or continue to drift. A three-type reality-infrastructure typology, comprising provenance systems, accountability structures, and legitimised contestation, was associated with substantially different outcomes. These case-level patterns are interpretive rather than statistical: across our cases, strong-infrastructure organisations recorded far fewer epistemic drift incidents (zero to one over

18 months) than weak-infrastructure organisations using comparable AI tools (six to eight). The findings contribute to debates on sociomaterial organising, the micro-foundations of institutionalisation, and the conditions under which professional judgment can be exercised under algorithmic co-production.

Keywords: *epistemic governance; reality infrastructure; confidence-decoupling; algorithmic co-production; professional services*

Introduction

When professionals produce knowledge, organisations have traditionally relied on two safeguards: formal controls (review, verification, and sign-off) and expert judgment (the trained eye of an experienced practitioner). Generative artificial intelligence (AI) weakens both at once. It produces content with no traceable author or reasoning chain to verify, and it generates the very surface signals, fluency, coherence, and professional formatting, that experts have long read as marks of quality. This paper asks how organisations keep their knowledge credible when neither safeguard works as it once did.

Sociomaterial perspectives suggest that the structure–agency dichotomy is mediated by technology: infrastructure shapes practice by becoming invisible (Star & Ruhleder, 1996), technology serves as an “occasion for structuring” that crystallises provisional practices into routines (Orlikowski, 2000), and algorithmic systems reshape how organisational knowledge is produced and evaluated (Kellogg et al., 2020; Leonardi, 2011). The conventional levers, tighter controls or more capable individuals, each prove insufficient on their own. This paper develops a theory of epistemic governance through infrastructuring to account for how organisations maintain credible knowledge not through individual virtue or structural control alone, but through the deliberate design of reality infrastructure: background systems that make epistemic status visible and actionable.

The study’s most striking observation previews the theoretical stakes. In organisations lacking governance infrastructure, we observed a confidence-decoupling phenomenon: a self-reinforcing pattern in which AI-enhanced surface quality raised

organisational confidence while objective epistemic quality declined. This connects to March's (1991) exploitation trap and Lindsley et al.'s (1995) efficacy-performance spirals, but specifies the infrastructural conditions under which organisations become trapped in self-reinforcing quality decline. The paradox, that organisations felt most confident precisely when their knowledge was least reliable, was largely invisible to the organisations experiencing it and became detectable only through the comparative design.

Through comparative case analysis of six professional services organisations (127 semi-structured interviews, 43 documents, and 22 observation sessions over 11 months), this study develops a grounded theory of epistemic governance through infrastructuring. The findings make four contributions. First, they develop the concept of reality infrastructure, the configuration of provenance systems, accountability structures, and contestability mechanisms through which organisations govern epistemic quality, and suggest that infrastructure strength depends on leadership intentionality rather than firm size or resources. Second, they offer a three-path process model showing how leaders' framing of epistemic shocks is associated with whether organisations build infrastructure or continue to drift, contributing to work on organisational sensemaking (Weick, 1995). Third, they describe the confidence-decoupling phenomenon, specifying conditions under which algorithmic augmentation may produce self-reinforcing epistemic failure. Fourth, they connect sociomaterial organising and professional jurisdiction theory (Abbott, 1988) by showing how AI may enable a self-inflicted form of jurisdictional erosion that is difficult for the affected profession to detect.

Theoretical Foundations

The Structure–Agency Impasse in Knowledge Governance

Structural accounts of knowledge governance emphasise how formal procedures, hierarchical review, and verification systems constrain action to produce reliable outcomes (Adler & Borys, 1996; Scott, 2001). Agentic accounts emphasise how professional expertise, situated judgment, and epistemic virtue enable practitioners to navigate uncertainty (Bandura, 2001; Emirbayer & Mische, 1998). Structuration theory

sought to transcend this dichotomy by proposing that structure and agency are mutually constitutive (Giddens, 1984), but its abstract recursiveness has proven difficult to operationalise in empirical research on knowledge quality (Barley & Tolbert, 1997).

AI co-production exposes the limits of both positions. Structural controls presuppose that the object of verification carries markers of its epistemic origins, an author, a reasoning chain, a source; AI-generated content carries none of these. Agentic accounts presuppose that professionals can assess quality through surface cues developed over years of practice; AI produces precisely those cues without corresponding epistemic grounding. The question becomes: what organisational mechanism governs knowledge quality when both structure and agency falter as conventionally theorised?

Sociomateriality, Infrastructure, and Algorithmic Organising

Infrastructure studies offer conceptual tools that leadership and knowledge-management literatures lack. Star and Ruhleder (1996) showed that infrastructure shapes practice by becoming invisible: once embedded, it recedes from awareness and its classificatory assumptions become taken for granted. Bowker and Star (2000) demonstrated that classification systems constitute what organisations treat as real. Orlikowski (2000) theorised technology as a catalyst that accelerates the crystallisation of provisional practices into routines, and Barley (1986) showed that new technology triggers structural change by altering the conditions under which actors interact. Leonardi (2011) developed the concept of sociomaterial imbrication, the interlocking of human and material agencies.

Recent work on algorithmic management extends these insights. Kellogg et al. (2020) identify mechanisms through which algorithms reshape organisational control; Möhlmann et al. (2021) examine how algorithmic management evolves with experience; and Dell'Acqua et al. (2023) show that AI augmentation produces heterogeneous effects on professional performance, with quality declining on tasks beyond the “jagged technological frontier.” This literature has focused on how algorithms control workers rather than on how organisations govern the epistemic quality of algorithmically co-produced knowledge. The present study addresses that gap.

Reality Infrastructure: Conceptual Foundation

We introduce the concept of reality infrastructure to capture the configuration of background systems through which organisations govern what they treat as epistemically credible. What distinguishes reality infrastructure from adjacent constructs is its function: it does not support cognition, store content, or administer control, it adjudicates epistemic status, making operational the boundary between claims the organisation treats as adequately grounded and those it does not.

This distinction sharpens against three neighbouring terms. Epistemic infrastructure, drawn from social epistemology, refers to the cognitive and social conditions supporting knowledge production, too cognitive and individualistic to capture the material, systemic dimension we observe. Knowledge infrastructure, common in information systems, refers to systems for storing, retrieving, and distributing knowledge, too focused on knowledge management to capture the evaluative function of determining whether claims are adequately grounded. Governance infrastructure, in institutional theory, refers to the broader apparatus of organisational control, too generic to capture the specifically epistemic dimension. We term our construct reality infrastructure to emphasise that these systems do not merely store knowledge or facilitate its flow but actively constitute what the organisation treats as epistemically real (Bowker & Star, 2000), by embedding provenance visibility, accountability assignment, and legitimised contestation into the material systems through which knowledge is produced, evaluated, and authorised.

Leadership-as-Practice and Infrastructuring

The leadership-as-practice (LAP) tradition frames leadership as collective, distributed, and materially embedded (Denis et al., 2010; Raelin, 2011). Crevani and Endrissat (2016) distinguish leadership exercised through materials from leadership as material practice. This study extends LAP by conceptualising epistemic governance as leadership achieved through infrastructuring, the ongoing material work of designing, maintaining, and enforcing systems that structure how knowledge is produced, evaluated, and authorised.

Professional Jurisdiction and Epistemic Authority

Abbott's (1988) theory of professional jurisdictions establishes that professions maintain social position through control over expert knowledge. Goldberg's (2010, 2018) account of epistemic dependence networks establishes that individuals routinely rely on others' testimony, and that this dependence rests on supporting infrastructure. Sperber et al.'s (2010) framework of epistemic vigilance describes how people calibrate trust using source-competence cues. AI co-production disrupts these networks, creating an epistemic heuristic gap in which the signals professionals rely upon become informationally impoverished.

Research question. How do organisations in AI-augmented professional services construct and maintain epistemic governance practices, and what forms does reality infrastructure take across different organisational contexts?

Method

Research Design

The study employed constructivist grounded theory (Charmaz, 2006, 2014) for three reasons: epistemic governance in AI-augmented settings is emergent and under-theorised; constructivist grounded theory acknowledges the researcher's co-construction of meaning; and iterative cycling between collection and analysis enabled progressive theory refinement.

Researcher Positionality

The researcher's 20-year professional background in professional services shaped both access and interpretation. This insider status facilitated rapport but created a risk of over-identification with managerial perspectives. Three strategies addressed this: systematic reflexive memoing (217 memos), deliberate prioritisation of junior professional accounts where they contradicted senior narratives, and member checking with 18 participants across four organisations. Early memos documented a tendency to read governance absence as negligence; prioritising junior accounts reframed it as a structural condition ("the system doesn't support checking") rather than individual laziness. Because some senior participants might perform competence in the presence

of a perceived peer, senior self-reports were triangulated with observational data, which frequently contradicted them, and probed through specific recent examples rather than general descriptions.

Organisational Recruitment and Access

Organisations were recruited through the researcher's professional networks in the sector. Access was negotiated at partner or practice-leader level, with written organisational consent and individual informed consent for each interview; participation was voluntary and confidential, and organisational and participant identities were pseudonymised. The researcher's sector background was material to access: partners at two firms explicitly noted that they were willing to discuss the "messy reality" of governance weaknesses because they perceived the researcher as understanding the commercial pressures they operated under, which enabled unusually candid accounts.

Case Selection and Theoretical Sampling

Professional services constitute a critical case for theory-building on epistemic governance: such firms sell credible judgment as their core product, exhibit high epistemic dependence on co-produced knowledge, have high exposure to generative AI, and produce observable failure outcomes when epistemic quality declines (Empson et al., 2015). Six organisations were selected through purposive and then theoretical sampling. Two maximum-variation cases anchored the study at the outset (StratConsult, a large firm with advanced AI use and a strong governance reputation; TechAdvisory, a mid-size firm with early AI use and minimal governance). As infrastructure configuration emerged as the key differentiating dimension, four further cases were added theoretically: GlobalLegal and RegionalLaw to develop and cross-validate an intermediate "partial" type; BoutiqueStrategy to test whether strong infrastructure could exist in a small firm (it could, disconfirming a scale assumption); and SpecialistAdvisory as a second weak-infrastructure case to test whether weak governance was a size effect (it was not). The sampling logic and the contribution of each case are documented in the theoretical sampling decision log in the online supplement, and case characteristics are summarised in Table 1.

Assessing Infrastructure Strength

Infrastructure type was not assigned from a pre-existing maturity instrument. At selection, organisations were sampled using observable proxies, stated AI-governance policies, reputation, and the presence or absence of provenance practices, to ensure variation. The formal classification into strong, partial, and weak infrastructure emerged from the analysis and was applied along the three dimensions that constitute reality infrastructure: provenance (systematic tagging versus opaque tracking), accountability (named owners at each stage versus diffuse, circular deferral), and contestability (structural protection of challenge versus suppression). Strong-infrastructure organisations exhibited all three in coherent combination; weak-infrastructure organisations exhibited none systematically; partial cases applied them inconsistently (for example, only on major pitches), producing the governance-leakage pattern described below. Classification was corroborated through member checking, in which all 18 participating informants confirmed their organisation's placement.

Defining Epistemic Drift Incidents

We define an epistemic drift incident as a discrete instance in which a weakly grounded or unverified AI-originated claim moved beyond provisional use to acquire institutional standing: a claim that, having entered without adequate provenance or verification, was subsequently reused, incorporated into a durable organisational artefact (a template, knowledge base, or client-facing deliverable), and treated as established fact, the cascade stage of the four-stage drift process (introduction, amplification, institutionalisation, cascade; see the coding tree in the online supplement).

Candidate incidents were identified by triangulating three converging sources: claim-level coding of proposal documents, participant accounts of specific failures elicited through the interview protocols, and organisational records of client-surfaced errors. To make counts comparable across organisations of differing size, incidents were referenced to the unit of the client engagement rather than the firm, and counted over an 18-month frame that combined the 11-month observation period with each organisation's retrospective records spanning the surrounding period. Classification was

undertaken by the researcher; each organisation's placement within the infrastructure typology was corroborated through member checking with 18 participants across four organisations.

Data Collection

In total, 127 semi-structured interviews were conducted at three hierarchical levels: senior leaders and partners (n = 38), mid-level managers and reviewers (n = 45), and junior professionals (n = 44). Interviews lasted 45–75 minutes, were audio-recorded with consent, and followed three role-specific protocols adapted iteratively as categories emerged (see the online supplement). Interview data were supplemented by document analysis (43 organisational artefacts) and direct observation (14 proposal-preparation sessions and 8 quality-review meetings).

Analytical Procedure

Analysis proceeded through three interconnected stages. Initial line-by-line coding generated 847 codes. Focused coding synthesised these into 34 categories. Theoretical coding identified relationships among 12 core categories organised around the central category of epistemic governance as infrastructural practice. Theoretical saturation was reached at approximately 110 interviews (no core category acquired new properties after Interview 94); 17 further interviews were completed for cross-case representation. The online supplement provides exemplar coding chains, the full coding tree, and the theoretical saturation matrix.

Findings

The central finding is that epistemic governance in AI-augmented professional services operates as an infrastructural practice, accomplished through the design of background systems that structure how knowledge is produced, evaluated, and authorised. Three core practices emerged, each a distinct dimension of governance through infrastructuring.

Structuring Knowledge Visibility

The most consistently observed practice was the enforcement of provenance disclosure at the point of judgment. At StratConsult, the managing partner introduced a three-tier provenance system tagging content as “human-originated,” “AI-assisted with human verification,” or “AI-generated pending review,” piloting it on his own engagements before firm-wide adoption:

“Before the tags, I was reviewing everything at the same level of detail. Now I know where to focus. The AI-drafted sections get a different kind of read.”

(SC-SA-04, Senior Associate)

BoutiqueStrategy’s founding partner developed a confidence-marker system assigning green/amber/red indicators by evidentiary strength, and personally reviewed every red-flagged claim in the firm’s first six months of AI use, modelling the behaviour the system was meant to institutionalise:

“If I skip the reds, my team will too. They’re watching what I actually do, not what the policy says.” (BS-P-01, Founding Partner)

By contrast, weak-infrastructure organisations had no systematic visibility into AI involvement. Origin blindness prevailed: once AI-generated content was edited, even the editing professional often could not identify which components were AI-produced (“After I edit the AI output, even I can’t tell which parts were mine and which were the model’s,” TA-JA-07, Junior Analyst).

Reconfiguring Review Authority

Strong-infrastructure organisations structurally legitimised challenge from junior professionals. At BoutiqueStrategy, an anonymous “red flag” system allowed anyone to question claims without attribution, and the founding partner publicly celebrated challenges, visibly protecting challengers:

“The first time a junior flagged something a partner had drafted, he thanked her in the team meeting. After that, people flagged things.” (BS-SA-02, Senior Associate)

This practice generated 3.8 challenges per proposal in strong-infrastructure organisations, compared with 0.8 in weak ones. The nature of challenges also differed: strong-infrastructure organisations saw substantive epistemic challenges targeting evidentiary quality, whereas weak-infrastructure challenges targeted formatting and style.

Enforcing Accountability Boundaries

The third practice assigned named epistemic ownership to defined knowledge claims:

“When my initials are on a section, I read it differently. I check the sources. I question whether the AI’s numbers actually add up.” (SC-P-06, Partner)

Verification rates under named ownership were 2.3 times higher than under diffuse ownership. Weak-infrastructure organisations exhibited circular accountability deferral, each professional assuming someone else had verified the content: “I assumed the analyst checked the numbers. He assumed I would. The manager assumed someone had” (GL-P-01, Partner). Configurations and outcomes across the three types are summarised in Table 2.

Infrastructure Emergence: A Three-Path Process Model

Cross-case analysis revealed three pathways through which infrastructure emerged or failed to emerge, with the critical branching variable being how leaders framed epistemic challenges (see Figure 1).

Path A: Reactive sensemaking. At StratConsult, infrastructure emerged after an epistemic shock, a client board identified unsupported claims in a strategy presentation. The managing partner framed this as a systemic problem requiring an infrastructural response rather than an individual failure requiring discipline. This systemic framing led to investment in provenance systems, accountability structures, and legitimised challenge. Drift incidents fell to zero to one over the following 18 months.

Path B: Anticipatory design. At BoutiqueStrategy, the founding partner built governance infrastructure from inception, before any failure occurred, reflecting prior experience at a larger firm where governance failures had been costly.

BoutiqueStrategy recorded zero drift incidents, achieved without the triggering failure that catalysed StratConsult.

Path C: Failure through individual-deficit framing. At TechAdvisory, a similar shock produced a different response. Leadership framed the failure as individual negligence and invested in training rather than infrastructure (“We sent the team on an AI prompt-engineering course. The problem was that people weren’t using the tools properly,” TA-P-01). The training addressed individual skills while leaving structural conditions unchanged; seven further drift incidents occurred over 18 months.

The Confidence-Decoupling Phenomenon

The most theoretically significant observation concerned the relationship between infrastructure and organisational confidence. In weak-infrastructure organisations, a four-stage self-reinforcing pattern emerged. AI-generated content increased the surface quality of outputs, which appeared more polished, comprehensive, and professionally formatted. This surface improvement raised organisational confidence in knowledge quality. Rising confidence reduced the perceived need for verification, lowering epistemic vigilance. Reduced vigilance allowed actual quality to decline further, increasing dependence on AI without commensurate governance.

The pattern was vivid at TechAdvisory, where it surfaced only through the comparative design. Senior participants described growing assurance, even as observational and document evidence pointed the other way:

“I’m feeling more confident with AI than I was six months ago. The output quality keeps improving.” (TA-SA-03, Senior Associate)

This account was recorded in the same period in which tracked claims at the firm were increasingly entering proposals without verification, and in which a client subsequently challenged a market figure that no one could trace to a primary source. Subjective confidence and objective quality were moving in opposite directions, and the divergence was itself the phenomenon, not a measurement complication. The pattern connects to Lindsley et al.’s (1995) efficacy-performance spirals at the organisational level and to March’s (1991) exploitation trap: success in AI-driven efficiency suppressed

exploration of governance alternatives. Critically, weak infrastructure did not merely fail to catch errors; it generated false confidence signals that accelerated quality decline.

Constructivist grounded theory generates theory rather than testing it. We therefore offer the following not as a verified finding but as a testable proposition that distils the cross-case pattern and invites confirmatory study:

Proposition 1. Under AI augmentation, organisational epistemic confidence will tend to decouple from objective knowledge quality when (a) surface fluency cues are high, (b) provenance opacity is high, and (c) verification infrastructure is weak, creating a self-reinforcing cycle that resists correction through the feedback mechanisms ordinarily available to the organisation.

Discussion

Infrastructuring as a Mode of Organisational Action

These findings suggest that infrastructuring, the deliberate design of background systems governing epistemic quality, operates as a mode of organisational action that offers a way of reframing the structure–agency impasse in knowledge governance. Infrastructure neither substitutes for professional judgment (the structural solution) nor relies on individual expertise alone (the agentic solution). Instead, it creates the material conditions under which professional judgment can be exercised: provenance systems make epistemic status visible so professionals can discriminate; accountability structures activate anticipatory verification; and contestability mechanisms provide the social conditions for substantive challenge. Read through the typology of provenance, accountability, and contestability, this helps specify, in this empirical setting, a mechanism that structuration theory (Giddens, 1984) left abstract: the configuration of infrastructure shapes which dimensions of agency are activated, without determining the content of action.

This extends sociomaterial perspectives on organising (Leonardi, 2011; Orlikowski, 2000) by identifying a form of material practice specifically concerned with epistemic quality. Whereas prior work on algorithmic management (Kellogg et al., 2020) examines how algorithms control workers, the present findings show how organisations can

design infrastructure that enables rather than constrains professional agency under algorithmic conditions.

The Confidence-Decoupling Contribution

Confidence-decoupling describes a mechanism of organisational failure under algorithmic augmentation. Unlike traditional quality failures, which produce visible signals that trigger correction, confidence-decoupling operates beneath the threshold of organisational awareness because the surface markers of quality improve even as the epistemic foundations deteriorate. This extends March's (1991) exploitation trap by specifying the infrastructural conditions that produce the trap under AI, and connects to Dell'Acqua et al.'s (2023) finding of heterogeneous AI-performance effects by suggesting an organisational mechanism through which negative effects may accumulate undetected.

Extending Professional Jurisdiction Theory

The findings suggest that AI augmentation may enable a self-inflicted form of jurisdictional erosion (Abbott, 1988) that is difficult for the affected profession to detect, operating through the confidence-decoupling mechanism. Unlike traditional jurisdictional threats from competing professions or technological displacement, this erosion originates within the profession's own knowledge-production practices: organisations can weaken the epistemic foundation of their authority without noticing, until a client or external party surfaces the failure.

Engagement With Practice

Although the study is observational rather than interventionist, it was designed to inform practice, and its action orientation operated by feeding findings back to participants in a usable form. Preliminary findings were shared with 18 participants across four firms in structured member-checking sessions during months 9–10; each session presented the emerging typology, the three core practices, and the confidence-decoupling pattern, and invited participants to confirm or contest their organisation's placement within the typology. Participants reported that the study's vocabulary, particularly "origin blindness," "confidence-decoupling," and "epistemic ownership," gave

them language for dynamics they had experienced but could not previously articulate, a framing that several described as directly usable in their own governance conversations.

Practical Implications for Leaders

Consistent with the journal's action-research orientation, five implications follow for organisational leaders adopting AI in knowledge work.

1. **Tag provenance before scaling AI, not after.** A simple three-tier tag (human-originated / AI-assisted-verified / AI-generated-unverified) redirected review effort to where it was needed. Retrofitting such visibility once workflows have hardened is substantially more costly than building it from the outset.
2. **Put a name on each claim, and have leaders model it.** Named epistemic ownership was associated with markedly higher verification. Assigning a named owner per section before circulation, with leaders visibly performing the same scrutiny, frames verification as expected rather than punitive.
3. **Make challenge cheap and safe.** An anonymous flag, combined with a leader publicly protecting the first junior challenger, shifted challenges from style toward substance.
4. **Frame the first failure as systemic, not personal.** The three-path model is the practical heart of the study: leaders who treated an epistemic shock as a system problem built durable infrastructure, whereas those who treated it as individual negligence bought training and continued to drift. Briefing leaders on this fork before a shock arrives may shape the response.
5. **Expect to justify success in risk terms.** Effective governance is self-effacing: the incidents it prevents never occur. Leaders should therefore measure and report avoided risk rather than efficiency gains alone.

Limitations and Future Directions

The professional services context bounds generalisability to knowledge-intensive settings. The 11-month observation period captured infrastructure emergence but not long-term sustainability, and the single-researcher design carries interpretive limitations

despite systematic reflexivity. The case-level drift comparisons are interpretive rather than statistical. Future research should test whether the infrastructure typology applies beyond professional services, examine the temporal dynamics of confidence-decoupling through longitudinal measurement, and investigate how infrastructure design interacts with evolving AI-system characteristics.

Conclusion

When AI co-produces organisational knowledge, maintaining credibility is less a matter of individual vigilance or structural control than of deliberate infrastructural design. Organisations that recognised this built reality infrastructure, provenance visibility, accountability structures, and legitimised challenge, that enabled professional judgment to operate under conditions of epistemic opacity. Organisations that framed epistemic challenges as individual deficits invested in training that left structural conditions unchanged, entering a confidence-decoupling cycle in which rising confidence masked declining quality. Epistemic governance through infrastructuring reframes the structure–agency impasse by creating the material conditions under which professional agency can be exercised without being determined, a form of organising that becomes increasingly foundational as algorithmic co-production reshapes knowledge work.

Data and Methods Transparency

This study was conducted as part of doctoral research. A methodological supplement, interview protocols, coding tree, power-quote table, theoretical saturation matrix, and sampling decision log are available at <https://osf.io/fw5kj> (<https://doi.org/10.17605/OSF.IO/FW5KJ>). Interview transcripts cannot be shared owing to participant confidentiality.

References

- Abbott, A. (1988). *The system of professions: An essay on the division of expert labor*. University of Chicago Press.
- Adler, P. S., & Borys, B. (1996). Two types of bureaucracy: Enabling and coercive. *Administrative Science Quarterly*, 41(1), 61–89.

- Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973–989.
- Bandura, A. (2001). Social cognitive theory: An agentic perspective. *Annual Review of Psychology*, 52, 1–26.
- Barley, S. R. (1986). Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Administrative Science Quarterly*, 31(1), 78–108.
- Barley, S. R., & Tolbert, P. S. (1997). Institutionalization and structuration: Studying the links between action and institution. *Organization Studies*, 18(1), 93–117.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). Association for Computing Machinery.
- Bowker, G. C., & Star, S. L. (2000). *Sorting things out: Classification and its consequences*. MIT Press.
- Charmaz, K. (2006). *Constructing grounded theory*. Sage.
- Charmaz, K. (2014). *Constructing grounded theory* (2nd ed.). Sage.
- Crevani, L., & Endrissat, N. (2016). Mapping the leadership-as-practice terrain. In J. A. Raelin (Ed.), *Leadership-as-practice: Theory and application* (pp. 21–49). Routledge.
- Dell’Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayner, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality (Working Paper No. 24-013). Harvard Business School.
- Denis, J.-L., Langley, A., & Rouleau, L. (2010). The practice of leadership in the messy world of organizations. *Leadership*, 6(1), 67–88.

- Emirbayer, M., & Mische, A. (1998). What is agency? *American Journal of Sociology*, 103(4), 962–1023.
- Empson, L., Muzio, D., Broschak, J. P., & Hinings, B. (Eds.). (2015). *The Oxford handbook of professional service firms*. Oxford University Press.
- Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30(4), 681–694.
- Giddens, A. (1984). *The constitution of society: Outline of the theory of structuration*. University of California Press.
- Goldberg, S. C. (2010). *Relying on others: An essay in epistemology*. Oxford University Press.
- Goldberg, S. C. (2018). *To the best of our knowledge: Social expectations and epistemic normativity*. Oxford University Press.
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410.
- Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2024). Generative AI and the disruption of expertise: The effects of AI writing assistance on expert knowledge networks. *Organization Science*, 35(1), 279–302.
- Leonardi, P. M. (2011). When flexible routines meet flexible technologies: Affordance, constraint, and the imbrication of human and material agencies. *MIS Quarterly*, 35(1), 147–167.
- Lindsley, D. H., Brass, D. J., & Thomas, J. B. (1995). Efficacy-performance spirals: A multilevel perspective. *Academy of Management Review*, 20(3), 645–678.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87.
- Möhlmann, M., Zalmanson, L., Henfridsson, O., & Gregory, R. W. (2021). Algorithmic management of work on online labor platforms: When matching meets control. *MIS Quarterly*, 45(4), 1999–2022.
- Orlikowski, W. J. (2000). Using technology and constituting structures: A practice lens for studying technology in organizations. *Organization Science*, 11(4), 404–428.

- Raelin, J. A. (2011). From leadership-as-practice to leaderful practice. *Leadership*, 7(2), 195–211.
- Scott, W. R. (2001). *Institutions and organizations* (2nd ed.). Sage.
- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origg, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359–393.
- Star, S. L., & Ruhleder, K. (1996). Steps toward an ecology of infrastructure: Design and access for large information spaces. *Information Systems Research*, 7(1), 111–134.
- Weick, K. E. (1995). *Sensemaking in organizations*. Sage.

Table 1

Participating Organisations: Case Characteristics

StratConsult	Consulting	>50,000	26	Strong	Yes	Systemic	0–1
BoutiqueStrategy	Strategy	<200	18	Strong	No	Anticipatory	0
GlobalLegal	Legal / M&A	8,000	22	Partial	Yes	Mixed	3
RegionalLaw	Commercial law	~1,200	20	Partial	Minor	Technical	2–4
TechAdvisory	Tech advisory	500–2,000	23	Weak	Yes	Individual	7
SpecialistAdvisory	Tech boutique	<200	18	Weak	Yes	Individual	6–8

Note. Infrastructure type emerged from analysis. Drift = epistemic drift incidents over 18 months.

Table 2

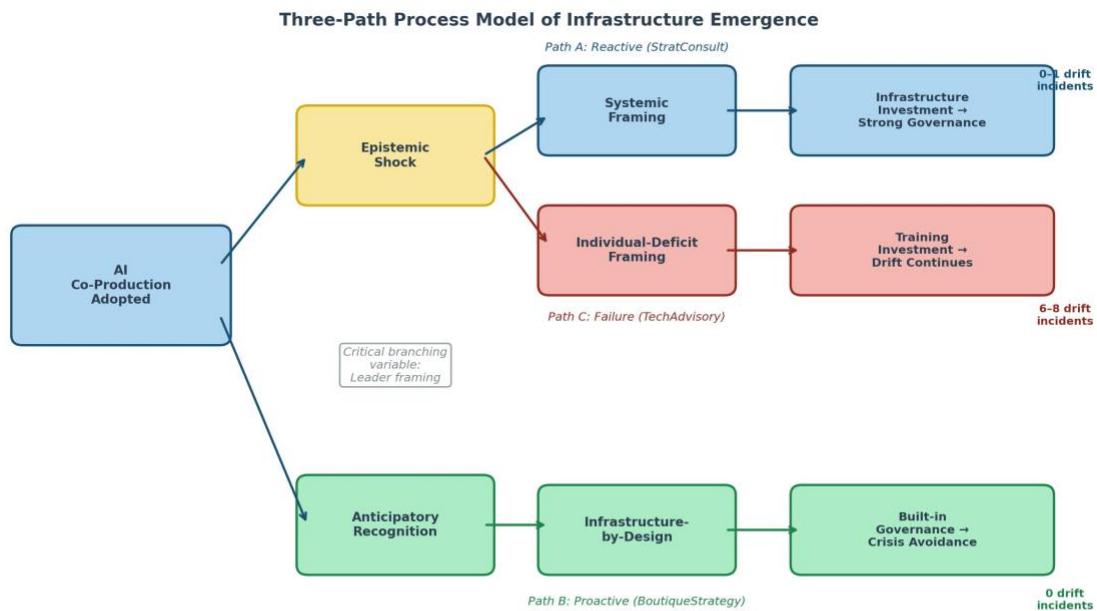
Reality-Infrastructure Typology: Configurations and Outcomes

Provenance	Systematic tagging	Some workflows; inconsistent	Opaque; no tracking	Targeted review
Accountability	Named owners at each stage	Major pitches only	Diffuse; circular deferral	2.3× verification
Contestability	Structural protection	Person-dependent	Suppressed	3.8 vs 0.8 challenges
Drift (18 mo)	0–1	2–4	6–8 (incl. failures)	Strong: near-zero
Confidence	Calibrated (stable)	Moderate gap	Decoupled (widening)	Most dangerous

Note. Cell entries summarise the modal configuration observed for each infrastructure type.

Figure 1

Three-Path Process Model of Infrastructure Emergence



Note. The critical branching variable is leaders' framing of the epistemic shock. Systemic framing (Path A) and anticipatory recognition (Path B) lead to infrastructure investment and near-zero drift; individual-deficit framing (Path C) leads to training investment and continued drift.