INNOVATION & RESEARCH CAUCUS

DIFFUSION OF TRUST IN AI IN PROFESSIONAL SERVICE FIRMS

How Trust Accelerate Technology Adoption in Professional Services Firms?

IRC Report No: 081

REPORT PREPARED BY

Francisco Trincado-Munoz
Oxford Brookes University

Hannah Tornow
Oxford Brookes University

Hilary Smyth-Allen
Oxford Brookes University

Sebastian Herbert-Fuchs
Oxford Brookes University

Delivered with ESRC and Innovate UK

Authors

The core members of the research team for this project were as follows:

- Dr Francisco Trincado-Munoz (PI) – Oxford Brookes University
- Dr Hannah I. Tornow - Oxford Brookes University
- Hilary Smyth-Allen – Oxford Brookes University
- Dr Sebastian Herbert Fuchs - Oxford Brookes University

This document relates to IRC Project FFTwo014 - Diffusion of Trust in Technology: Understanding how Tech Adoption spread in Professional Service Firms

Acknowledgements

This work was supported by Economic and Social Research Council (ESRC) grant ES/X010759/1 to the Innovation and Research Caucus (IRC).

We are very grateful to the project sponsors at UK Research & Innovation (UKRI) for their input into this research. The interpretations and opinions within this report are those of the authors and may not reflect the policy positions of UKRI or its constituent councils. We would also like to acknowledge and appreciate the efforts of the IRC Project Administration Team for their support in preparing this report.

About the Innovation and Research Caucus

The Innovation and Research Caucus supports the use of robust evidence and insights in UKRI's strategies and investments, as well as undertaking a co-produced programme of research. Our members are leading academics from across the social sciences, other disciplines and sectors, who are engaged in different aspects of innovation and research systems. We connect academic experts, UKRI, IUK and the (ESRC), by providing research insights to inform policy and practice. Professor Tim Vorley and Professor Stephen Roper are Co-Directors. The IRC is funded by UKRI via the ESRC and IUK, grant number ES/X010759/1. The support of the funders is acknowledged. The views expressed in this piece are those of the authors and do not necessarily represent those of the funders.

Contact

You are also welcome to email us if you have any questions about this report or the work of the IRC generally: info@ircaucus.ac.uk

Cite as: Trincado-Munoz, F., Tornow, H.I., Smyth-Allen, H., and Herbert-Fuchs, S.H., (2026). *Diffusion of Trust in AI in Professional Service Firms*. Oxford, UK: Innovation and Research Caucus

CONTENTS

Executive Summary

Professional service firms (PSFs) are under growing pressure to adopt generative AI (GenAI), yet adoption often stalls, not because the technology fails, but because people hesitate to trust it. This report examines how trust in GenAI spreads through a mid-sized UK law firm, and how that "contagion" can be harnessed to accelerate responsible adoption. Drawing on a longitudinal, mixed-method case study (including 81 interviews, 78 surveys, ethnographic observation, and a social network analysis across two points in 2025), we trace how trust in AI is built, shared, and sometimes withheld in everyday professional practice.

Our central finding is that trust in GenAI is contagious but fragile, and that it is produced together with distrust rather than simply replacing it. Professionals extend trust task by task. That is, they rely on GenAI for low-risk work while firmly withholding it from high-stakes tasks (e.g., court preparation or client advice), where human oversight remains non-negotiable. This "efficiency–risk paradox" sits alongside deeper tensions around professional judgement, client-data duties, career progression, and the fear of de-skilling.

Crucially, trust does not trickle down the organisational hierarchy. The network analysis shows that it spreads sideways, through peers, and is brokered by unexpected actors, notably junior, technology-fluent staff whose informal influence far exceeds their formal rank. A colleague sitting alongside a hesitant user, or a shared time-saving success, builds trust more than any top-down mandate. At the same time, distrust travels through its own channels: risk encounters, the opacity of AI systems, and suspicion directed at "big tech."

For firms, the practical lesson is to lead by enabling rather than mandating. Surface and support peer champions and junior brokers, protect time to experiment and to coach, make safeguards visible, and build the oversight and verification norms that keep high-stakes use safe. Firms are currently under-prepared for the supervisory shift AI brings, as partners can no longer assume that work handed up has travelled through familiar, trusted routes.

For policymakers, the priorities are to translate principles-based regulation into practical, sector-specific safeguarding and oversight templates through professional bodies; to embed AI literacy and human-oversight competence in accreditation and continuing professional development; and to fund peer-led, SME-focused diffusion and cross-firm communities of practice, targeting the mechanisms of trust rather than the volume of tools procured.

Adoption succeeds not when AI is imposed, but when firms make trust safe to give and easy to pass on. Trust in GenAI is built by people, through people and supporting the relationships through which it travels is the most effective lever available to firms and policymakers alike.

1. Context

Many legal firms today have advanced in AI adoption. According to the Solicitors Regulation Authority (2023), around three-quarters of the largest firms already use AI, and over 60% are exploring GenAI. However, legal firms often struggle to engage users in the adoption process and reduce resistance to change. Resistance to change arises from stringent regulatory frameworks, complex hierarchies, and a billable-hour-focused business model. Finding mechanisms to accelerate adoption is an urgent need for firms to fully leverage technology for innovation. Trust in technology has been identified as an important precursor to its adoption. Some researchers have explained that, by means of trust, users will be more confident using technology, unlocking barriers to change. Thus, this project aims to understand the following question: How can trust in AI accelerate the successful adoption of technology?

The UK's leadership in professional services depends on integrating technologies into traditional practices. In the legal sector, UK firms are investing in AI for growth and competitiveness. For small and mid-size firms, AI (specifically GenAI) offers opportunities to expand services. Specifically, UKRI has prioritised AI adoption to foster an innovation-led economy (e.g., Next Generation Professional and Financial Services program or the Technology Adoption Accelerators).

Legal professionals are cautious about adopting AI because of the high-risk nature of their work and the industry's dominant business model. A key challenge for PSFs has been accelerating AI adoption, with trust identified as a catalyst. Trust can enable open discussions about AI risks, building confidence that AI will support rather than threaten professional roles. Establishing trust in AI is key, and not only suppliers play a role in sustaining it; peers are essential in promoting trust and reducing perceptions of AI risk.

Our project examines how trust in AI spreads through a legal firm, creating a “contagion effect” that can drive AI adoption across the firm. We aim to identify the factors and variables which influence this contagion. Our findings aim to guide innovators and policymakers on the conditions that support adoption, offering evidence on how to build trust in AI.

While much focus has been on building trustworthy technology, less attention has been given to fostering trust in its adoption. In legal firms, a strong professional culture, driven by client focus and billable hours, often hinders tech adoption. A report by the Law Society of Scotland (2024) identified people and culture as key challenges in AI adoption. Indeed, technology adoption should not be top-down, where managerial boards and innovation teams mandate technology use, but should engage the individuals who will use it in their legal work from the outset. Our project addresses this gap by examining how trust in AI develops among frontline professional users through peer interactions, building on insights from our previous work at the Technology in Professional Services (TiPS) accelerator.

1.1. The Professional Service Firms sector and the AI adoption challenge

In the UK, for the professional service firm sector to be leading is dependent on integrating digital technologies, including AI. In the legal sector, UK firms are investing in AI for growth and competitiveness. For small and mid-size firms, AI offers opportunities to expand services and stay competitive. We draw on insights from previous research from the TiPS (Technology in Professional Services) accelerator, an accelerator for small and medium-sized law and accountancy firms to adopt technology. As part of the TiPS research project, we learned that trust is at the heart of successful technology adoption. Leveraging these findings, we go in depth into one case study to understand how trust in AI diffuses and what factors influence the diffusion of trust. The organisation we studied is a mid-sized law firm in the UK committed to adopting technology, specifically AI, through Microsoft Copilot.

1.2. Aim and Objectives

This project sets out to understand how trust in AI can be fostered and diffused inside a professional firm, and how that contagion of trust can be encouraged to support the responsible adoption of the technology. Rather than treating adoption as something switched on by senior management (top-down imposition), we look at how trust in GenAI is built, shared, and sometimes withheld in the everyday interactions of the people who actually use it. To do this, we pursue four objectives:

- Understand the individual, relational, and organisational factors that determine whether trust in AI spreads or stalls within a firm.
- Identify the key tensions that underpin how trust in AI is built and transmitted.
- Examine how trust develops as professionals put AI to work in high-stakes, regulated tasks.
- Translate these insights into practical guidance for firms, technology providers, and policymakers who want to accelerate adoption without undermining professional standards.

2. What do we know?

Before turning to our own work, it is worth setting out what existing research already tells us about trust in technology and AI. We will do this in three sections: (1) what generative AI is and why it matters for professional firms; (2) why trust is central to whether people adopt a new technology; and (3) how trust diffuses in organisations.

2.1. Generative Artificial Intelligence (GenAI): Why do we care?

GenAI refers to a class of tools, most visibly the large language models behind ChatGPT, Claude, and Microsoft Copilot, that can produce text, summaries, code, and analysis in response to everyday language prompts. It has been described as a "general purpose technology", whose effects are not confined to one task or sector but ripple across the whole economy, much as electricity or the internet did before it (Eloundou et al., 2024). Hence, its potential is still unknown.

The interest in GenAI is not hypothetical. Multiple studies have shown meaningful productivity gains, yet also substantial risks. More recently, GenAI has evolved to produce Agentic AI. For instance, customer-support agents resolve issues faster and more successfully (Brynjolfsson et al., 2023). Likewise, management consultants complete knowledge tasks with AI faster and produce higher-quality output. For small and mid-sized firms in particular, AI offers a route to expand services and stay competitive, and UK policy has actively encouraged this shift through programmes aimed at building an innovation-led economy.

Yet the qualities that make GenAI useful also make it risky in a legal setting. Some studies have warned of a "jagged frontier", where the tool quietly performs worse on tasks that sit just outside its competence, precisely where an unwary user is most likely to be caught out (Dell'Acqua et al., 2026). Models can produce fluent but incorrect answers, raise questions about client confidentiality, and sit awkwardly with a profession built on accountability and precedent. The Law Society of Scotland (2024) identified people and culture, not the technology itself, as the central challenge in adoption. This is why trust, rather than raw capability, so often becomes the decisive factor in whether AI is actually used.

2.2. The Role of Trust

As new technologies emerge and professionals engage more with technology, trust begins reshaping how they feel about this (Messerschmidt & Hinz, 2013). Trust is commonly defined as a willingness to be vulnerable to another party based on positive expectations of their behaviour (Mayer, Davis and Schoorman, 1995). When the "other party" is a technology rather than a person, users weigh strikingly similar concerns: Is the system competent, reliable, and working in my interest? (McKnight et al., 2011). Trust in the context of GenAI becomes a difficult construct to understand, as neither a relational nor a technological approach seems to fit.

Research on automation shows that trust allows people to rely on a system they cannot fully monitor, and that both too little and too much trust are problematic. Under-trust leads people to ignore genuinely useful tools, while over-trust leads them to accept flawed output uncritically (Lee and See, 2004). Trust in AI is especially fragile because these systems are opaque and their reasoning is difficult to inspect, so users have little to go on beyond the results they see and the opinions of those around them (Glikson and Woolley, 2020).

Decades of technology-adoption research confirm that people adopt tools they find useful and easy to use (Davis, 1989; Venkatesh et al., 2003). But in high-stakes professional work, those judgements are filtered through trust. Where the cost of an error is a professional's reputation or a client's case, confidence that AI will support rather than expose them is a precondition for use. Trust also enables the frank conversations about risk on which safe adoption depends.

When trust is absent, professionals will tend to disengage, reverting to established routines even when the technology offers demonstrable advantages (Glikson & Woolley, 2020). Conversely, when trust is sustained, it lowers the psychological cost of experimentation, facilitates learning-by-doing, and enables the iterative engagement through which technologies become embedded in organisational practice (Li et al., 2008). Trust is therefore not merely a condition for initial adoption but a dynamic resource that sustains and deepens technology use over time, making it central to any account of how organisations convert technological investment into innovation capacity. However, in high-stakes professional contexts, the question of who and what professionals trust, and how those trust dispositions travel through organisational networks, is a missing mechanism in the adoption literature (Goto, 2023). This situation is exacerbated by the disruptive nature of GenAI, as well as the inexplicable features of LLMs that make users uncomfortable (Trincado-Munoz et al., 2025).

2.3. The Diffusion of Trust

How does a new technology move through an organisation? Rogers' (2003) classic account of the diffusion of innovations shows that adoption spreads unevenly through social channels, with respected "opinion leaders" and change agents shaping whether others follow. Adoption is, in large part, a social process rather than an individual decision.

Crucially, it is not only the technology that spreads; trust in it does, too. People take their cues about whether to rely on something from the colleagues around them, a phenomenon often described as social contagion (Burt, 1987). The structure of relationships matters. For example, close ties transmit confidence and reassurance, while "weak ties" and brokers who bridge otherwise separate groups carry new information and ideas across a firm (Granovetter, 1973; Burt, 1992).

Who/What influences who is not random. For example, people tend to trust and take advice from those similar to them in role, seniority, or background (a tendency known as homophily: McPherson, Smith-Lovin and Cook, 2001). This suggests that hierarchy, profession, and demographics will shape the paths along which trust in AI travels, and that the same colleague who reassures one person may hold no sway over another.

Taken together, this research tells us that trust is central to adoption and that it spreads socially, but it says little about how trust in AI specifically diffuses inside a professional firm, or how it interacts with distrust as people begin to use the technology in confidence. That is the gap our study addresses.

3. Methods

We have chosen a law firm in the UK with a strategic plan to adopt GenAI, specifically Microsoft Copilot. The firm was contacted through a program funded by the UK government which supported PSFs to accelerate technology adoption (TiPS). Access was negotiated with a gatekeeper, who is leading the firm's digital transformation. The firm is based in Scotland with offices across the country. A group of partners (the board) and a managing partner oversee day-to-day operations. This group is committed to implementing GenAI as part of their growth strategy in the next few years.

Supporting this plan, the firm created a transformation department in charge of the implementation and adoption of GenAI, as well as the cultural transformation tied to this. The main initiatives of this group can be summarised here:

- Increasing the company's workers' general technology literacy – by implementing the LTC4 Framework.
- Human-centred design – initiatives with people at the centre, considering human-in-the-loop and engaging people.
- Safe experimentation – implementation of a sandbox for experimenting, testing, and iterating GenAI.

To increase engagement, the company invited employees to become GenAI adoption champions, increased communication through different channels (especially implementing a mobile tech hub, where experts from the transformation department will go to the different offices to approach employees, solve technical questions, and encourage training), and created some measures to track Return-on-investment (ROI).

During 2025, drawing on a longitudinal mixed-method study, we collected three different types of data. We collected interviews and surveys at two points to understand changes in trust in AI.

- Interviews: Participants were approached independently of the firm and the gatekeeper, following ethical guidelines. In total, we interviewed 45 respondents at two time points (45 in T1 and 36 in T2).
- Surveys: We also collected survey data: 78 surveys (45 in T1 and 33 in T2). For the surveys conducted, respondents were asked to name colleagues who first come to mind when they think about (1) who is an expert on AI (learning network), and (2) who they will go to for advice on AI (advice network). In addition, we used two scales to measure trust in AI (15 items) and readiness (4 items). We also obtained information on professionals' gender and role in the firm.
- Ethnographic observation: We participated in 8 participant ethnographic observations. We conducted non-participant observations to capture naturally occurring discussions among employee groups, associated challenges and concerns, and other AI-related developments within the company. Additionally, we observed the intervention with the travelling tech hub across various sites.

In addition, we complemented our data collection with available documentation (flyers, reports, presentations). It is important to note that a number of participants left the organisation during the two data collection points. The second wave was also scheduled shortly before the Christmas period, during which increased sickness absence and end-of-year workload pressures limited participants availability.

While the qualitative data was analysed using Gioia's methodology (2013) (see Appendix 1). Quantitative data, including network observations, was analysed using social network analysis. We triangulated both type of data with observation notes and documentation, so we could understand how trust emerges and diffuses, and actors relate to each other in the process of adoption.

4. Findings

4.1. Network Analysis

We focused the analysis on respondents-only directed networks (i.e., a tie A→B means A named B as an expert / an advisor). While wave 1 (T1) was collected in June 2025, wave 2 (T2) was collected in December 2025. In the first wave, we used responses from 35 of the 45 participants; we removed 10 responses because they did not complete the questionnaire or we could not locate their nominations in the organisation. From these, only 28 also responded to the survey in wave 2. We anonymised each participant. Table 1 presents the network characteristics for each time point.

Table 1. Network Characteristics for Expertise and Advise at both time points

network	nodes	ties	density	Centralisation (indg)	Components	Top In-degree	Betweenness (Avg)
Expertise T1	35	47	0.039	0.353	3	LQ(13), AMO(12), JMC(5)	0.0039
Expertise T2	13	10	0.064	0.562	3	LMY(7), CC(1), IB(1)	0
Advise T1	33	36	0.034	0.223	3	LQ(8), AMO(7), JMC(6)	0.0021
Advise T2	16	12	0.05	0.373	4	LMY(6), IB(2), LW(2)	0

Table 1 shows different centrality measures. (1) Density shows the number of existing ties from the potential ties between nodes. For example, a density of 0.039 for the Expertise Network at T1 means that only 3.9% of the possible (100%) ties appear in the network. (2) Centrality (In-degree), shows how many times a person has been named (i.e., "who is sought out" (have more prestige)). Finally, betweenness represents how often a node lies on shortest paths between others (i.e., has the role of "brokerage"). When this value is zero, it means the network has components that cannot be joined.

A first analysis shows that both networks shrank sharply between waves (Experts 47→10 ties, Advice 36→12) because the two dominant hubs (LQ and AMO) left the firm. Yet the surviving networks became more centralised, not less. In-degree centralisation rose (Experts 0.353→0.562; Advice 0.223→0.373) as advice-seeking re-concentrated on a single successor (LMY). Betweenness is low throughout and effectively zero at T2: these are radial "go-to" structures (everyone points straight at a hub), not chains, so almost no one brokers between others. At T1 the only meaningful brokers were the hubs themselves (Experts betweenness: JMC 0.054, LQ 0.036, AMO 0.029).

Below we present the visualisations of the networks. Node size in every figure represents the in-degree centrality (number of times a person was named). Layout is shared across T1 and T2, so each person keeps the same position in both panels and the structural change is visible directly.

4.1.1 Expertise Network

At T1 the expert network is a dense, multi-hub star (see Figure 1). LQ (in-degree 13), AMO (12) and JMC (5) absorb most nominations, with JMC also the top broker (betweenness 0.054) because they link the LQ and AMO clusters. By T2 the structure has collapsed to a single dominant hub, LMY (in-degree 7), plus a few disconnected dyads (CC–IM, GMA–LW). Density rises (0.039→0.064) only because the node count fell faster than the tie count; substantively the network is smaller and more dependent on one person. The rise in centralisation (0.35→0.56) means expertise-seeking is now a **single point of dependency**, a resilience risk if LMY were also to leave. In both networks we see that fee people (LQ, AMO, JMC, LMY) concentrate the expertise and advise roles.

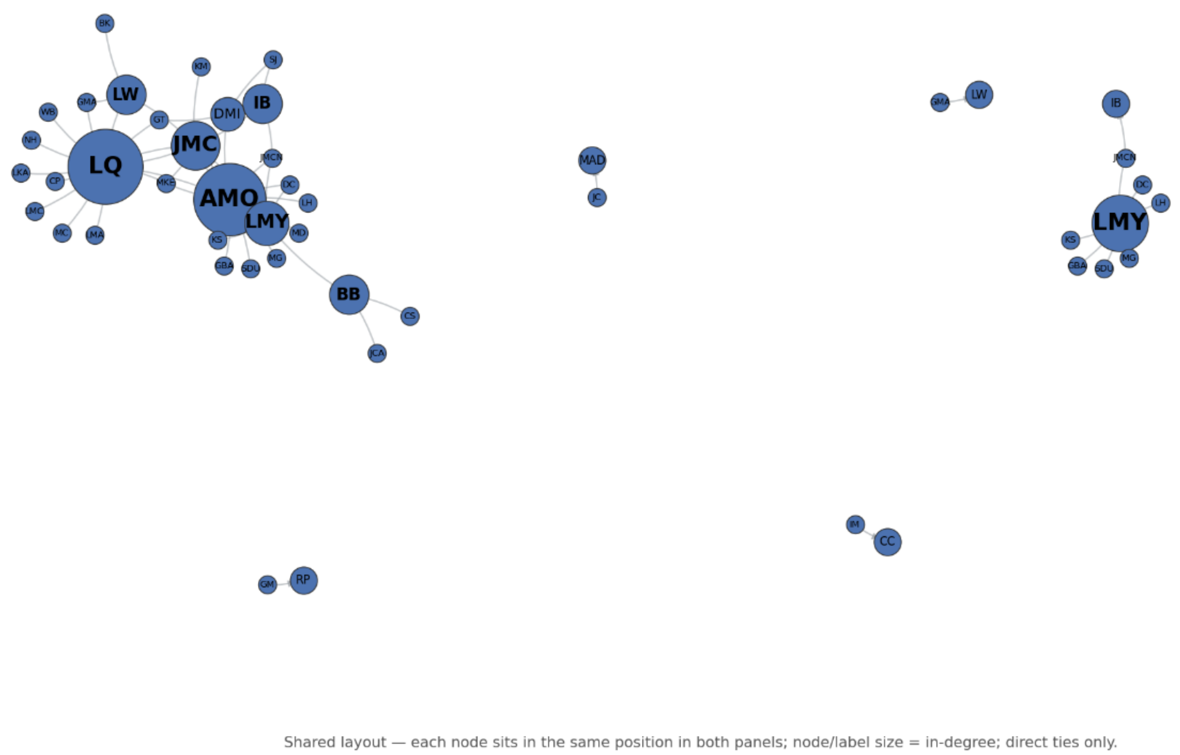

Figure 1. Expert Network (T1 - left vs T2 - right)

At T1 the top experts are mixed-gender, LQ (male) and AMO/JMC (female) are all central, and mean in-degree is essentially equal (Female 1.32 vs Male 1.38) (see Figure 2). At T2 centrality tilts strongly female, mean in-degree Female 1.00 vs Male 0.25, and the sole hub (LMY) is female. This is a by-product of who remained at the firm rather than a gendered preference in nominations; the departing hubs happened to include the most-named male (LQ).

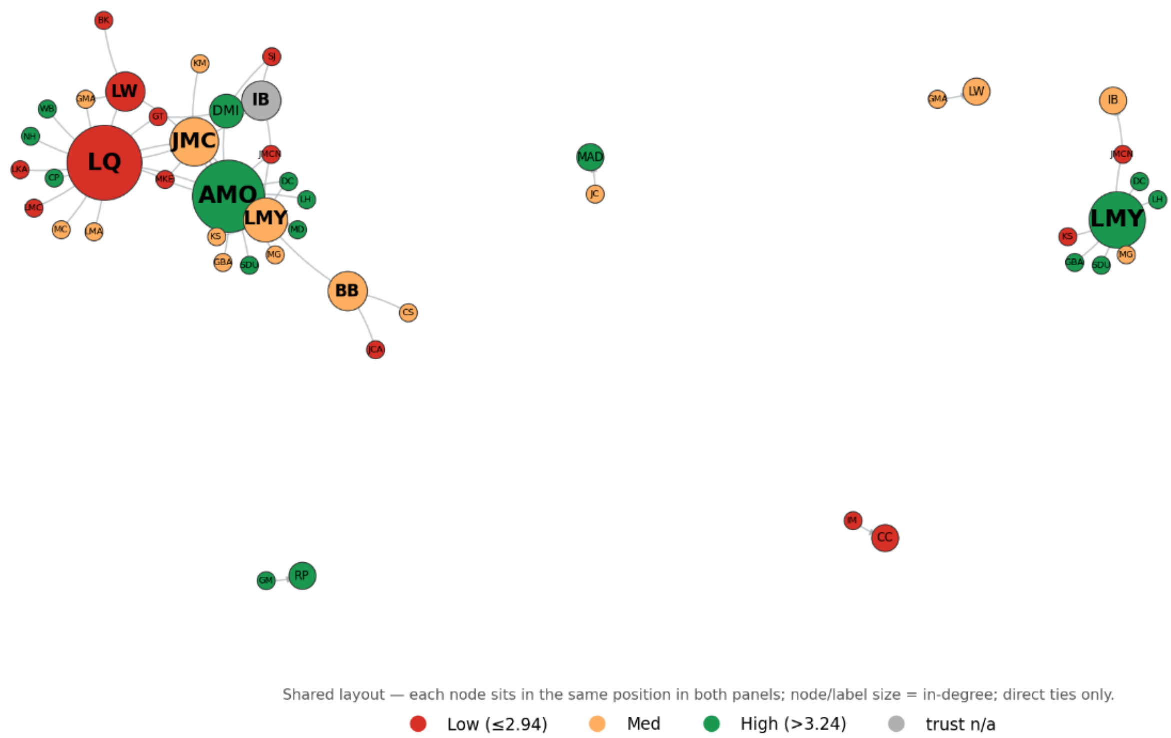

Figure 2. Expert Network (T1 - left vs T2 - right) by Gender.

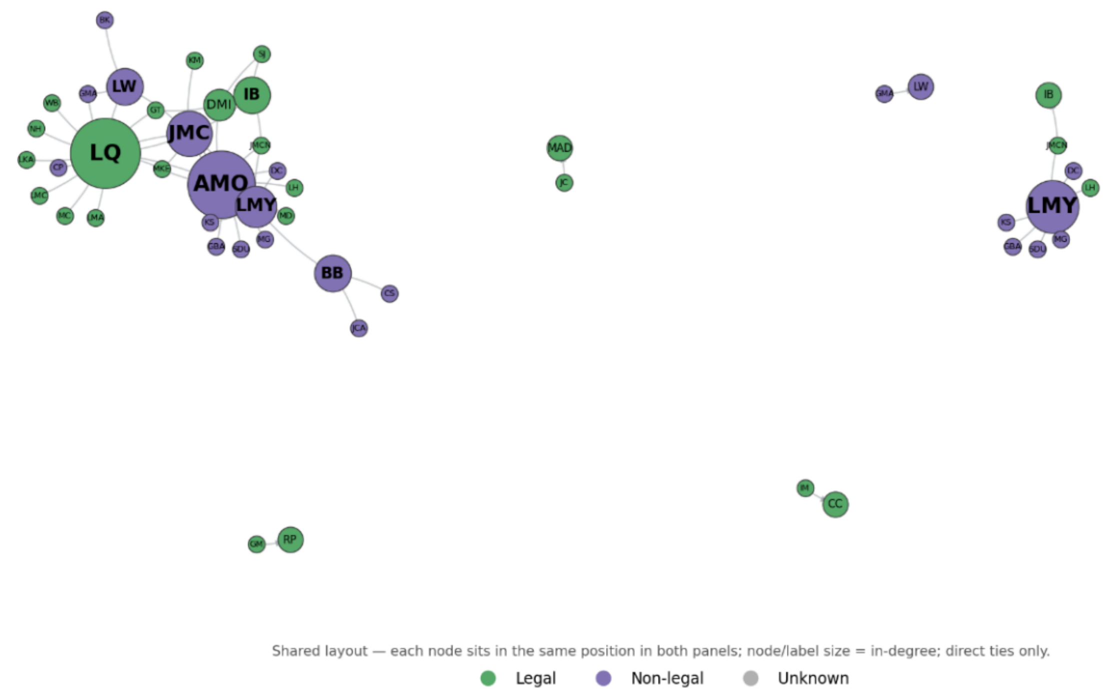

Figure 3. Expert Network (T1 - left vs T2 - right) by Role.

Non-legal staff (IT/AI support, marketing, operations) are consistently more central as experts than legal staff (mean in-degree Non-legal 1.80 vs Legal 1.00 at T1, and 1.00 vs 0.40 at T2) (see Figure 3). This is intuitive, people seek technical AI expertise from the firm's technical/support functions. The pattern is stable across waves even though the specific individuals changed, indicating a structural (role-based) source of expertise rather than reliance on particular people. LMY, the T2 hub, is a non-legal role, continuing the pattern.

The most interesting cross-wave finding. At T1, the most-sought experts were *low*-trust (mean in-degree Low 1.60 > High 1.33 > Med 1.00), driven by LQ (low trust, in-degree 13) (see Figure 4). At T2, this flips. High-trust nodes are now most central (mean in-degree High 1.40 vs Low 0.25), because the successor hub LMY is High-trust. Thus, the firm's go-to expert shifted from someone comparatively *skeptical* of the AI tool to someone comparatively *trusting* of it. If expertise carries attitudes, this could gradually raise trust among those who consult the hub, a plausible diffusion mechanism.

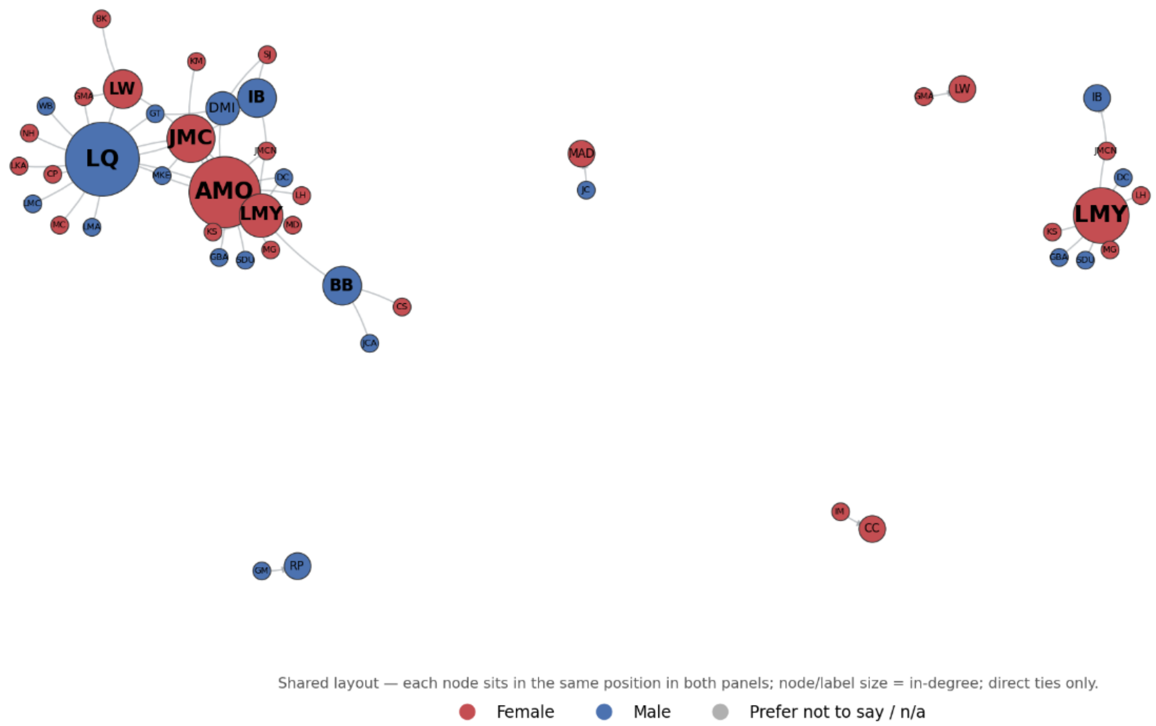

Figure 4. Expert Network (T1 - left vs T2 - right) by Trust Level.

4.1.2 Advise Network

The advice network mirrors the expert network but is slightly less centralized at both waves (centralization 0.22 vs 0.35 at T1) (see Figure 5). At T1 four people share the load, LQ (in-degree 8), AMO (7), JMC (6), BB, and brokerage is again concentrated in JMC/AMO/LQ (betweenness 0.020/0.018/0.014). At T2 advice re-centralizes on LMY (in-degree 6), with small satellites (IB, LW at in-degree 2). Components rise to 4. That is, the advice network is now fragmented into one LMY-centred core plus isolated pairs.

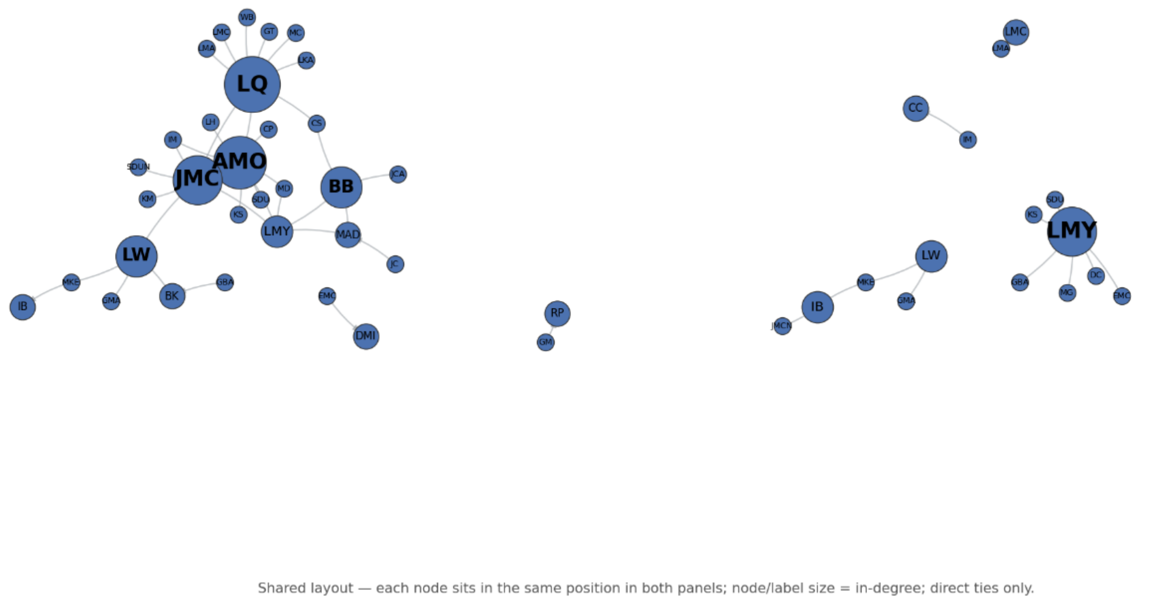

Figure 5. Advise Network (T1 - left vs T2 - right)

Women are somewhat more central as advisors in both waves, and the gap widens over time (see Figure 6). Mean in-degree Female 1.24 vs Male 0.94 at T1, and 1.13 vs 0.38 at T2. As with experts, this reflects hub survival (LMY) rather than a gendered nomination bias, but the practical effect is that advice-giving at T2 rests largely on female staff.

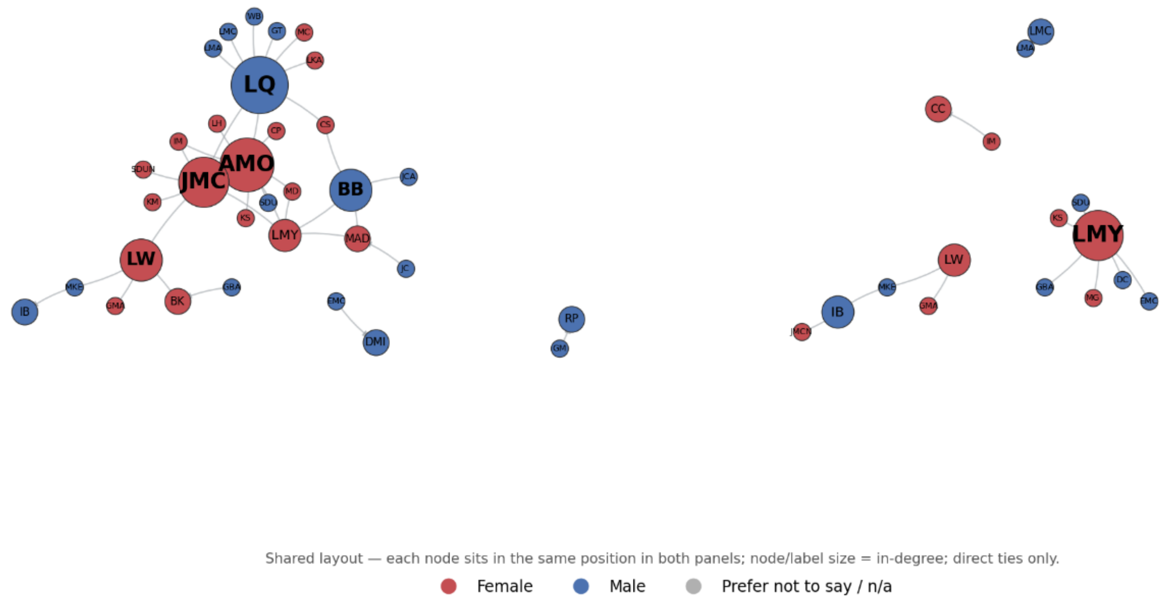

Figure 6. Advise Network (T1 - left vs T2 - right) by Gender

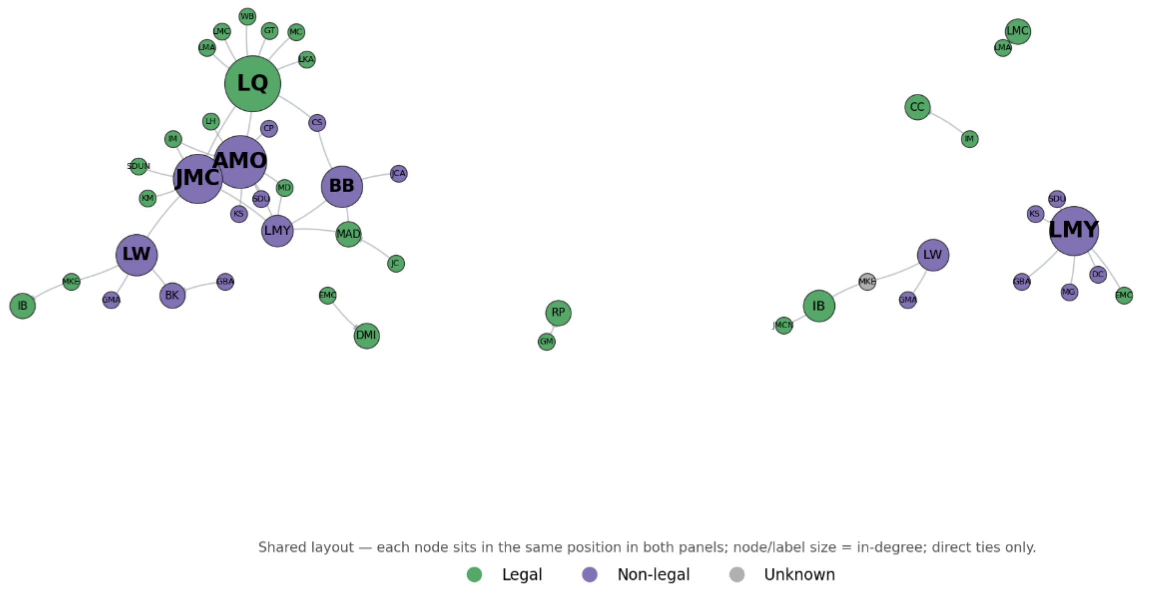

Figure 7. Advise Network (T1 - left vs T2 - right) by Role.

The role pattern is even sharper for advice than for expertise (see Figure 7). At T1 Non-legal staff have mean in-degree 1.85 vs 0.60 for Legal. That is, advice on AI flows overwhelmingly toward the technical/support functions. At T2 the gap narrows (Non-legal 1.00 vs Legal 0.57) as the network thins, but Non-legal remains the more consulted group and LMY (non-legal) is the hub. Legal staff are largely advice-seekers rather than advice-givers in this domain.

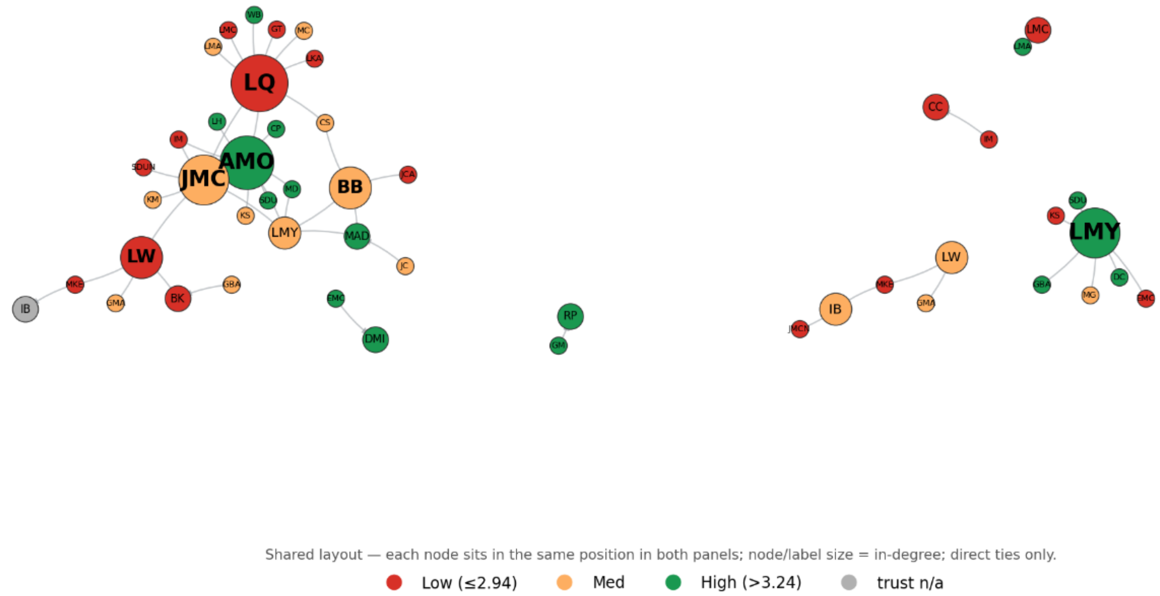

Figure 8. Advise Network (T1 - left vs T2 - right) by Trust Level.

The same trust flip seen for experts appears for advice (see Figure 8). At T1 the most-consulted advisors skew *Low*-trust (mean in-degree Low 1.30 > Med 1.09 > High 0.91), again driven by LQ. At T2 the ordering reverses to favour *High*-trust advisors (High 1.20 > Med 1.00 > Low

0.29), driven by LMY. Across both networks, then, the people the firm relied on for AI guidance moved from relatively risk-aware/skeptical (T1) to relatively trusting (T2).

4.1.3 Comparison built by characteristics

For further exploring the distribution on participants characteristics, we show two figures for Gender (Male/Female) and Role (Legal/Non-legal). In each figure, we show the average for the readiness, trust in technology, and the four dimensions of trust (Functionality, Predictability, Competence, and Risk Perception (Low-risk)). We place each measure side by side across waves (blurry = T1, solid = T2).

Overall, trust and readiness barely moved (General Trust 3.12 → 3.12; Readiness 3.72 → 3.71), but that stability hides two shifts. Perceived risk rose (Low-Risk fell 2.61 → 2.48) and Competence perceptions dropped (3.38 → 3.22), while Predictability improved (3.08 → 3.27), that is, people find the tool more predictable but see it as riskier and less capable than they did two months earlier. In more detail:

Figure 9 shows that men are consistently more ready and trusting than women in both waves, and the gap is stable. The clearest movement is women's Low-Risk score dropping (2.31 → 2.18), women perceive more risk at T2.

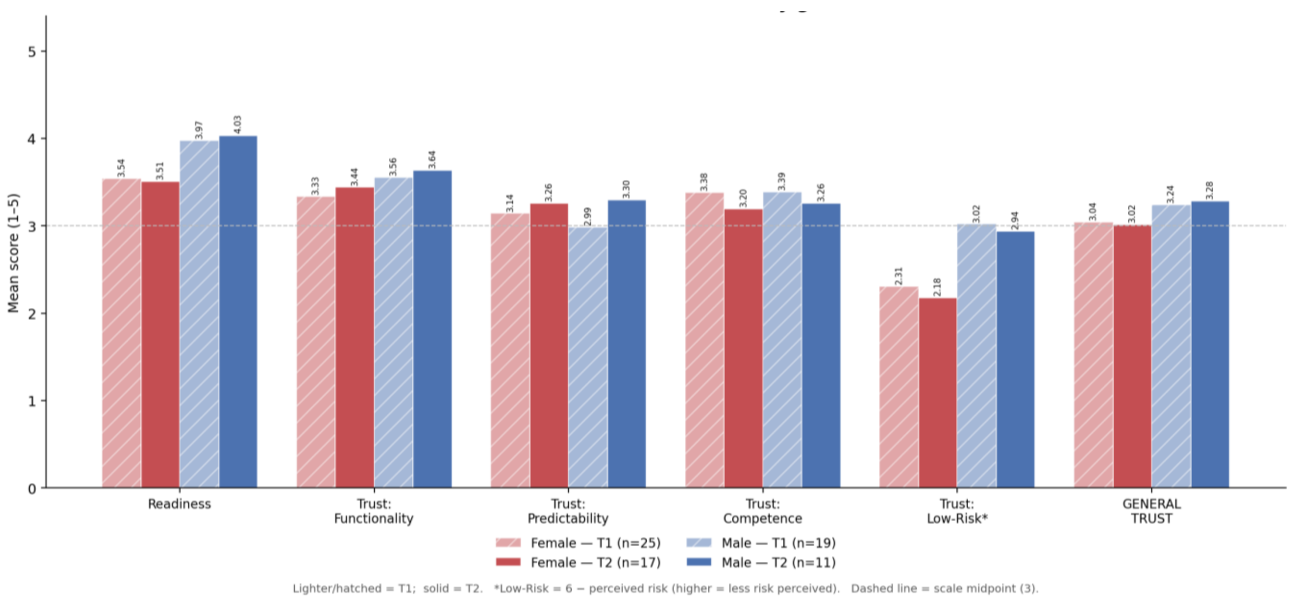

Figure 9. Trust and Readiness by Gender.

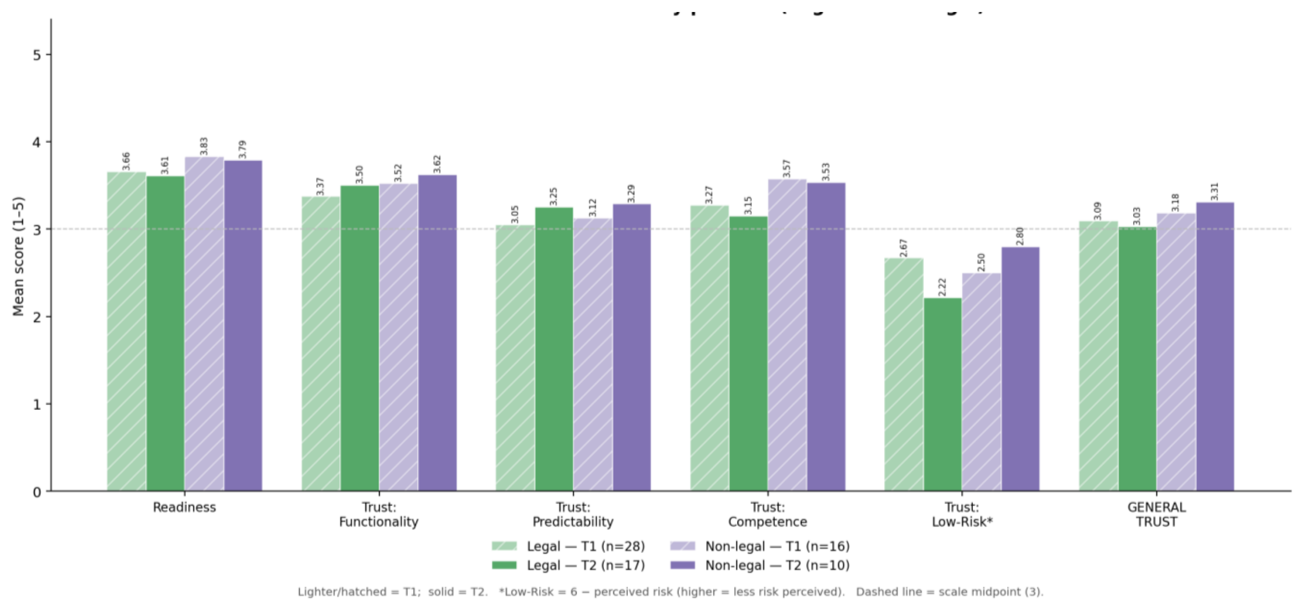

Figure 10. Trust and Readiness by Role.

Legal and Non-legal track closely on General Trust (see Figure 10), but Legal staff's Low-Risk score fell sharply (2.67 → 2.22). Thus, the legal group grew notably more risk-aware, whereas Non-legal moved the opposite way (2.50 → 2.80).

4.2. Insights from Interviews

The 81 interviews with professionals at the law firm reveal interesting insights regarding trust in AI, and how it diffuses across the organisation. First, there is a recognition that AI is inevitable and has gone beyond the initial “hype”. Second, trust grows with repeated use, yet time and skills are barriers that create hesitation (vulnerability). Third, we found that trust goes along with distrust, both of which diffuse across people in the organisation and become a safe combination for professional adoption of AI. Indeed, five different forms of vulnerability shape the tension between trust and distrust: institution, team, peers, colleagues, and third parties. Finally, we also discuss some potential barriers and enhancers of trust.

4.2.1 A point of no return: from ‘hype’ to inevitability

Participants no longer framed GenAI as a “hype”. Instead, adoption was widely experienced as inevitable, a shift already underway that the firm had little choice but to engage with. In general, there was an acknowledgement of a mandate from the leadership to explore AI and an obligation to adopt this to remain competitive. This sense of a ‘point of no return’ set the emotional backdrop against which trust was then negotiated.

4.2.1.1 From hype to inevitability

Rather than dismissing AI as over-hyped, participants described it as a permanent and accelerating feature of professional life, and framed engagement as the pragmatic response.

> *"I don't think it's a hype, I think…it's coming, whether you, like, we like it or not, I think it won't die down if it's just going to pick up pace, get better and better, so it's better to jump on the train and go with it."* – *Partner (T06_T1)*

> *"I think it's past the hype stage, and probably at the implementation stage now, but carefully."* - *Senior lawyer (T07_T1)*

4.2.1.2 Top-down pressure to experiment

This sense of inevitability was reinforced from above. Leadership signalling and the visibility of AI at external events created a felt pressure to be seen to be adopting it. Participants also recognise incentives to adopt AI.

> *"All about AI and how we're looking to use it within the firm. So that's, again, promoting AI, and quite often, if I go to any kind of conferences or presentations AI's kind of beginning to feature more and more, and definitely being driven, I would say, from the top."* – *Partner (T06_T1)*

4.2.2 Constraints shaping initial trust

While the direction of travel felt settled for professionals, day-to-day engagement was shaped by practical constraints. Three recurred more often in the respondents: (1) trust had to be earned through use; (2) skill and knowledge gaps create hesitation; and (3) access to fee-earners vs non-fee-earners for questions and advice.

4.2.2.1Trust grows with time and repeated use

Informants framed trust not as a starting condition but as something that would increase gradually, through familiarity and exposure. Ease of use and opportunities for testing help to reduce concerns and disperse vulnerability threats. Time is the answer for many to increase adoption.

> *"I think that's (trusting the technology) is going to come with time, and using it more. I will become more comfortable with it, and what it's saying, and what it's generating."* – *Business Support staff (T.11_T1)*

> *"Giving it time, then I'll become more confident in it."* — *Business support staff (T.11_T1)*

4.2.2.2 Skill and knowledge gaps create hesitation

A lack of confidence and, in particular, an absence of training left willing users hesitant to engage, and afraid of doing their work or unsure what the tool was even for. This is connected to the lack of training, familiarity with the tool. Furthermore, having the knowledge and skills to be able to understand the tool (AI literacy) and use this.

"I played about with it a little bit to say, for instance - 'we don't want to go there, can you suggest somewhere else and update the whole itinerary' - and played around with it. I think, from my point of view, I feel I need a bit more confidence in challenging it, because I worry that if I change something, I'll lose everything that I've already done." – Partner (T.06_T1)

"[I am] keen to use it, but I'm hesitant until I've added any training, so I've not had any training whatsoever yet, so I've not used it…until you've had that training, I suppose you don't know what you can use it for." – Partner (T.17_T2)

4.2.2.3 When time is a barrier: lawyers versus non-lawyers

Because fee-earners' time is billable and scarce, users evaluate whom they approached for guidance, often turning to support colleagues rather than busy lawyers, a pattern that shaped how knowledge and reassurance travelled. The consequence is that testing is always based on the technical aspects, and less on the legal expertise and the quality of the outputs that the technology provides. The culture of protecting the fee-earner remains a barrier for searching support when experimenting with the technology.

"I would… I wouldn't go to X, because I feel he's a fee earner, he would be far too busy to be dedicating time to… pointing me in the right direction, whereas Z (non-lawyer professional) would be my go-to person." – Business support staff (T.11_T1)

4.2.3 The paradox of trust diffusion: trust and distrust

These constraints on the initial experimentation with technology lead to a more substantial aspect. At the heart of the adoption of technology lies a paradox. That is, our findings show that professionals hold not only trust in GenAI, but a combination of trust and distrust that are constantly in tension. The two emerged together and fed one another: while professionals trust the technology, they also distrust it. In other words, scepticism persisted even among users who valued the tool, and reassurance rarely settled their concerns about the technology.

"I think, what keeps a lot of people away from it, they just don't believe, and they're a bit sceptical that it actually will do what they need it to do, especially until they've actually seen it and seen how it works" – Partner (T.05_T1)

"I don't find it 100 % reliable, the results are never exactly accurate (…) From a business point of view, it is not quite trustworthy yet" - Technical Support Staff (T.22_T2)

The paradox of trust and distrust was also sustained by a set of fears and dilemmas, each of which is examined below.

4.2.3.1 The “Efficiency–Risk” dilemma

Trust was always task-contingent. The same tool that was trusted for low-stakes work was firmly distrusted for high-stakes tasks; quite often, it was tolerated only as a starting point, **always** subject to human oversight. Professionals keep weighing risk, as they should, not only professional risk but also personal (e.g., reputation).

> *“For certain tasks, it can be trustworthy, because there's not so much risk with it, like with the blog post, but you wouldn't fully trust it to do the preparation for the court case by itself. No, absolutely not, on its own, but it's good to give you a starting point, because it does”* – *Senior lawyer (T.07_T1)*

4.2.3.2 ‘Bounded use’ of AI: navigating the compliance maze

Fear of doing something ‘wrong’ led to a bounded, heavily-checked mode of use of AI. Compliance anxieties also reframed supervision: the trust once extended to colleagues risked being displaced by a mistrust of AI-mediated work. Keeping a sight on AI outputs is the rule, avoiding over-reliance.

> *“That… that element of fear, I suppose, don't we? And the risk to make sure that …We aren't doing anything that's wrong, but I can't, I don't see actually any way around it, other than just making sure everybody is checking things very carefully before they rely on it.”* – *Partner (T.5_T1)*

> *“So, so long as, let's say, a note of advice to a client makes sense you can pretty much trust that they've got the law right, otherwise they wouldn't be employed here. But with the introduction of AI. I think that may actually make it more difficult to supervise, because you're taking that level of trust away, it is being replaced with mistrust of AI"* – *Partner (T.06_T1)*

4.2.3.3 ‘Judgement tension’: identity conflict

GenAI touched a nerve of professional identity. Participants insisted human judgement remained essential, stressed non-negotiable duties around client data, and worried that reliance on AI might erode the deep grounding that lets experienced lawyers sense when something “looks off”. There is also a tension on the expertise, and how knowledge is compared/contrasted to AI, and a new type of law professional emerges (lawtech lawyer).

> *"As more people who are used to AI get more senior. I have a slight worry that some of that older knowledge about how to go and test the base information and the source information might be slightly lost. I may be being I may be doing my junior colleagues a disservice. Maybe it won't, maybe because they've grown up with the technology, they'll know automatically that they have to check it, but it worries me a little bit that a whole generation who've always known Google and*

> *are very trusting of technology get into supervisory roles, where they might not be as questioning as perhaps my generation and older are."* – *Partner (T.29_T2)*

4.2.3.4 Fear of job replacement

Beneath the reassurances that the firm offers, it lies an anxiety about redundancy, a worry that, as the tool improves, clients and colleagues might feel they no longer need the professional at all. The fear of being replaced remains salient, and lawyers and non-lawyers feel this threat.

> *"A lot of people are scared that it would go away without the human input. Actually, we're not really needed if these sorts of things get more and bigger and bigger? And as people get more comfortable using them, do they feel there are things they could produce themselves? And they therefore don't need us, you know, to be doing things for them?"* - *(T.05_T1)*

4.2.3.5 Fear of loss of legal skills

A distinct fear concerned de-skilling. That is, the reliance on AI might make professionals faster without making them better, reducing the core competencies juniors are meant to develop. This creates hesitation that transfers to the relationship with junior lawyers.

> *"Junior lawyers should get used to actually not using it in the first instance, because they want to learn the core skill, and then you should more use it as, like a convenience (…) But I mean, I think even at my stage, you will probably end up loosing skills, because you are relying on AI to help you (…) It is beneficial for my time, but not for my development (…) It is helping you be faster, but is it helping you be better, like without the tool, do you still have the same skills? Will you still have the same capabilities, or are you replacing your capabilities with the tool"* - *Senior Lawyer (T.30_T2)*

4.2.4 Sources of vulnerability that affect trust and distrust

Trust and distrust did not sit only in individuals; they were transmitted and contested across several levels of the firm. Through the institution, teams, colleagues, peers, and third parties, individuals feel reassured or more hesitant about AI, consequently affecting their trust/distrust. Thus, these sources of vulnerability diffuse trust/distrust, where tensions emerge, and contagion happens.

4.2.4.1 Institution / organisation

At the broadest level, trust is underwritten by the firm itself. That is, knowing that internal systems have been examined (e.g., going through procurement), and providers have been evaluated, creates reassurance. Once firms give users permission to proceed, or, conversely, prevent them from engaging with a technology, these signal potential vulnerabilities to users. The institution and its leadership are the first to guarantee experimentation and further use.

> *"I'm pretty confident it's more trustworthy than it isn't, because I know the systems that we have, as far as I know, have all been checked through from our internal guys, for us to safely use it." - Business support staff (T.12_T1)*

4.2.4.2 Colleagues

The endorsement from colleagues, for example, by doing comparable work, is also a powerful driver for trust/distrust. A time-saving success story quickly translates into motivation to engage with the technology. Likewise, a bad experience will prevent/cut engagement. There seems to be a level of vicarious experience with the technology, a narrative exercise, where stories of failure or success navigate across people, so they serve as reassurance/warning for using the technology.

> *"Positive feedback from colleagues, if a colleague in the team is saying 'I used Copilot to do this', it could save time. You know that is great, that is the ultimate goal. And if someone is doing the same job as me and saves time, well, obviously, yeah, I want to do that too." - Senior Lawyer (T.14_T1)*

4.2.4.3 Teams

Beyond a single colleague, the team's overall viewpoint and experience also seem to have important consequences for trust/distrust. Indeed, teams acted as vehicles for spreading know-how, with workshops and joint projects deliberately used to circulate knowledge. Acting as a group seems to favour the members' trust/distrust and create a reference through which users feel validation to use the technology or prevent them from doing so.

> *"There are certain people who are working in teams that are doing workshops and projects with the IT team on AI things. I suppose that's what these workshops are now for, is to spread that knowledge around." - Partner (T.05_T1)*

4.2.4.4 Peers

Perhaps the most potent mechanism was direct peer reassurance. A peer sitting alongside a hesitant user and validating their approach converted anxiety into confidence. A peer's negative experience, especially in cases of client duty, became a clear warning of the potential faults of the technology.

> *"Seen a couple of bits that {they}s been doing, and I actually wasn't particularly comfortable that I was using it properly. So [X] actually came and sat with me at my desk and said. Show me what you're doing, and I showed them the different things I was doing. X went, no, you're absolutely using it. You're at the start of your journey, this is exactly what we'd be expecting you to do."- Business support staff (T.11_T1)*

4.2.4.5 Third parties

Third parties, such as social media, providers, regulators, and professional bodies, are also sources of vulnerability. In particular, distrust is fed by what professionals hear from these third parties. If the big-technology providers behind the tools cannot guarantee the security of personal and client data, engaging with the technology will be unlikely to happen. Indeed, the opposite- the silence of some related parties also speaks to hesitation (at some points, even validation – if they do not say anything, it probably means it is fine) or validation (they don't care).

> *"Question mark [to whether it is trustworthy]. I don't know. I don't know, I don't… I don't trust big tech…I'm always thinking someone's trying to steal my details somewhere." – Business support staff (T.09_T1)*

4.2.5 Additional barriers and enhancers of trust / distrust

Finally, the diffusion of trust was moderated by a set of barriers and enhancers. The most visible was a generational dynamic that cut both ways. Age and seniority slowed adoption, yet junior colleagues emerged as an unexpected catalyst.

4.2.5.1 Barrier: the generational gap

Inexperience with the technology was concentrated among more senior staff, which was a recurrent brake on adoption in the case organisation. Senior staff recognised a steeper learning curve and, quite often, a lack of time to engage with the technology because of their commitments. Older professionals tend to be linked to resistance to change; also, their level at the firm gives them more protagonism, so more junior staff observe and monitor their experience.

> *"Still early days, like, I'm still fairly kind of new to it all, I had never really… thought to use AI before for anything, so yeah, it's a bit all new to me. I'm not… Hugely great with technology." – Partner (T.6_T1)*

> *"I mean, the partners, they're all a bit older, on the older side, or, like, in the later stages of their careers, just because they're partners. Um, and they're usually the most reluctant to use any new technology, and they struggle the most to adopt it as well." – Junior Lawyer (T.33_T2)*

4.2.5.2 Enhancer: the role of junior lawyers

Juniors became trust brokers. Precisely because they were comfortable with technology, they were entrusted with AI, used it more readily, and, despite their formal rank, tend to be taken seriously as the 'go-to' people, reshaping how AI supervision works. There is also a recognition of the responsibility this represents, and the risks for this traditionally hierarchical profession (e.g., junior lawyers will not learn the fundamentals or become less able to think critically) where juniors take control of some knowledge, balancing power based on knowledge.

> *"Well, I suppose if you look at, my colleague for an example, is technically a trainee solicitor he's kind of, in terms of the hierarchy at the bottom, both of us are but because he's doing this stuff with AI and tech in that respect, he is kind of the go-to person for that, so people higher up in the hierarchy would probably have those discussions with him. And despite the fact that he's technically at the bottom of the ladder. He's not… that doesn't reflect in the importance that he holds in respect of things that he says about the tech and the AI, like, they're taking him seriously, and it's not dismissed just because he's a trainee." - Junior Lawyer (T.41_T2)*

4.2.5.3 Barrier: change in supervision and distrust

The flip side of juniors' new role was a supervisory strain. Partners could no longer assume that work handed up had been produced through familiar, trusted routes, so trust in the junior risked being replaced by mistrust of the AI behind them. There is constant (unnecessary) guessing about whether something has been produced with AI, and traditionally, cues are sought to confirm supervisors' suspicions or not.

> *"[If] a note of advice to a client makes sense, you can pretty much trust that they [Junior Lawyers] have got the law right; otherwise, they wouldn't be employed here. But with the introduction of AI. I think that may actually make it more difficult to supervise, because you're taking that level of trust away, it is being replaced with mistrust of AI" – Partner (T.06_T1)*

4.3 Final Figures

Engagement with the firm's AI initiative grew markedly over the course of the research and intervention. Membership of the AI special interest group rose from 53 to 79 members, an increase of 49.1%, indicating a widening base of active participation rather than passive awareness. Over the same period, the number of Microsoft Copilot users grew by 75%, and by the close of the intervention, 46 members of staff had completed AI training, with a further 17 registered to do so.

By the end of the research, significant changes occurred in the transformation team, pausing the initiatives in the firm. Despite this, Microsoft Copilot license usage reached approximately 90%, without a structured AI activity such as formal training and the firm's engagement initiatives. A large share of staff with a license are still using AI without support. This has been translated into a pause in the diffusion and adoption of AI throughout the organisation. Without social scaffolding (e.g., peer support, training, and interventions), existing licenses or users do not translate into broader confidence and adoption across the organisation.

5. Implications and Insights

Our findings, and the accompanying social network analysis, carry practical lessons for professional service firms and clear signals for policymakers seeking to accelerate responsible AI adoption. Because trust in AI is co-produced with distrust and travels through people rather than mandates, the actions that matter most are relational and organisational, not merely technical. We group them under three themes: putting safeguards in place, closing the oversight gap, and harnessing the peer- and junior-led channels through which trust actually spreads. Policy implications are explained below.

The most consequential reframing our findings invite is that trust and distrust in AI are not two ends of a single dial but two independent dimensions; they are orthogonal (Lewicki, McAllister and Bies, 1998). A professional can hold genuine confidence in a tool and pointed scepticism about it at the same time, and our data show exactly this: professionals who relied on AI for low-risk drafting still withheld it from binding work (for multiple reasons), and endorsement from a trusted peer rarely erased the residual doubt attached to "big tech" or the “black box” of AI. The practical implication is that interventions which successfully build trust do not thereby dissolve distrust. Firms that treat the two as an easy-to-influence factor, assuming that more reassurance automatically means less resistance, will likely misread how professionals feel about adopting the technology. The task is not to convert distrust into trust, but to manage two diffusion processes running in parallel.

Seen this way, the diffusion of trust and the diffusion of distrust travel through the same social network but along partly different channels, and each needs its own strategy. Trust spreads through peer reassurance, visible time-savings, and the junior staff that are more familiar with the technology but have legal knowledge, too; distrust spreads through risk encounters, opacity, and suspicion projected outward into third parties (e.g., suppliers) or harvested by third parties (e.g., media outlets). Because the two are orthogonal, a firm can and should amplify the first without attempting to suppress the second. Indeed, a measured level of distrust is precisely what sustains the oversight behaviour that makes high-stakes use safe (Lee & See, 2004). The implication is that firms should cultivate a “trusting distrust”. That is, **enough confidence for people to engage, enough doubt for them to check**. Thus, interventions, such as peer champions, should be mandated to transmit both the enthusiasm and the cautionary norms, rather than acting as one-sided advocates whose success would lead the firm into uncritical over-trust.

The policy implication follows directly: innovation support and professional guidance should aim not to maximise trust but to keep the two dimensions in productive tension. Funding the diffusion of confidence while embedding the disclosure, verification, and oversight practices that give distrust a constructive outlet. Adoption succeeds not when distrust is defeated, but when trust and distrust are allowed to do their different jobs at once.

5.1 Safeguarding measures

The trust–distrust paradox means that adoption stalls where safeguards are invisible. Confidence rises sharply when users know a tool has been inspected internally, so institutional assurance is one of the strongest trust signals. Firms should therefore make safeguards explicit and visible. For example, clear rules on client-data handling and consent before any information is entered into a model; "human-in-the-loop" checkpoints calibrated to task risk, so that low-stakes work can proceed freely while high-stakes outputs are always verified; and promoting a culture of safe experimentation, where tensions are resolved.

Policy implication. Smaller firms rarely have the resources to design such safeguards from scratch. Professional bodies and regulators should translate high-level principles into practical, sector-specific templates, model data-handling clauses, consent wording, and risk-assessment checklists.

5.2 Not prepared for AI oversight

Firms are markedly under-prepared for the supervisory shift that AI brings. Supervision is being quietly reconfigured. On the one hand, partners can no longer assume that work handed up has travelled through familiar, trusted routes, so trust in the junior risks being displaced by mistrust of the AI behind them. On the other hand, the role that junior lawyers are taking as tech-savvy users and more confident in AI is levelling up power, undermining the knowledge hierarchy and creating tension. Because AI is still opaque, oversight cannot rely on inspecting the process, but it must focus on verifying outputs and preserving the professional "grounding" that lets experienced lawyers sense when something is wrong. In law firms, oversight capacity is unevenly distributed, generally concentrated in a handful of experienced partners; their lack of supervision can leave parts of the firm exposed to reputational or even financial risks.

Four potential measures to fight this: (1) verification norms (routine cross-referencing of AI-assisted work); (2) disclosure of where and how AI has been used in a work product; (3) training that deliberately protects core law skills, so that efficiency gains from AI do not become a loss of competence; and (4) rethink the legal professional career, with space for a new type of lawyer who combines both technical and subject-specific knowledge.

Policy implication. Regulators should set clear expectations for AI oversight, competence, and disclosure; the SRA has already encouraged firms to supervise AI as they would a junior employee. Rethinking how continuing professional development should embed AI-literacy and human-oversight skills, give opportunities for lawyers to have alternative careers, and offer funding to help smaller firms build oversight capacity (and not merely acquire tools).

5.3 Harnessing peer- and junior-led trust

The clearest insight from the social network analysis is that trust in AI does not flow down the organisational chart, but it spreads sideways, through peers, and is brokered by unexpected actors. Influence over whether colleagues came to trust AI mapped poorly onto formal seniority. Instead, junior, technology-fluent staff occupied central positions, bridging the firm's IT and innovation function and frontline legal practice. Connection to peers, a colleague sitting alongside a hesitant user, or a time-saving success story passed on, were the strongest conduits of trust, and people took the advice most readily from those similar to them in role and background.

Firms should mobilise these channels deliberately rather than leave them to chance. Precisely, (1) identify and support the internal champions (in engaging ways) - distinguish champions from influencers, too; (2) create structured peer coaching and "sit-with-me" sessions; (3) protect fee-earners' time for knowledge-sharing but still engage them; and (4) legitimise reverse mentoring so that juniors' expertise is recognised without destabilising the hierarchy. Because trust accrues with repeated use, protected time to experiment is itself a trust-building intervention, not a cost.

Policy implication. Government-funded innovation programmes, such as the TiPS accelerator, should back bottom-up, peer-led diffusion rather than top-down technology mandates, and should fund cross-firm communities of practice so that smaller firms can share the brokers and champions they cannot resource alone. Innovation policy should target the mechanisms of trust diffusion, not the volume of tools procured.

6. Conclusions

Trust in AI is contagious, but fragile. It is produced together with distrust, calibrated task by task, and carried through the everyday relationships of the people who use the technology. The network analysis shows that this contagion runs along peer ties and is brokered by junior, technology-fluent staff whose informal influence far exceeds their formal rank. Adoption therefore succeeds not when it is mandated from the top, but when firms make trust safe to give and easy to pass on.

The firm's own experience during the study bears this out. Over the course of the research intervention, active membership of the AI special interest group grew, with Microsoft Copilot users and training increasing substantially (evidence that trust, once seeded through peer support and hands-on help, spreads). However, without sufficient support those gains cannot be sustained and adoption stalls. While implementing AI (getting more licences) is important, sustaining the adoption through a social scaffolding (by coaching, training, and forums where users discuss AI) is how trust in AI actually diffuses and supports sustained adoption. The learnings from this research lead to some implications for policy and practice.

6.1 Implications for practitioners

- Lead by enabling, not mandating: surface and support peer champions and junior brokers rather than issuing directives.

- Make safeguards visible: vetted tools, clear data-and-consent rules, and human checks scaled to the risk of each task.
- Protect time to experiment and to coach, recognising that trust grows with repeated, supported use.
- Build oversight and verification norms, and protect the core skills that make meaningful oversight possible.

6.2 Implications for policy

- Translate principles-based regulation into practical, sector-specific safeguarding and oversight templates.
- Embed AI literacy and human-oversight competence in professional accreditation and continuing professional development.
- Fund peer-led, SME-focused diffusion and cross-firm communities of practice, targeting the mechanisms of trust rather than tool procurement.
- Support oversight capacity in smaller firms, so that responsible adoption does not depend on scale.
- If the goal is a professional-services sector that adopts AI responsibly and at pace, the lever is trust, and trust is built by people, through people.

REFERENCES

Brynjolfsson, E., Li, D. and Raymond, L.R. (2023) Generative AI at work. NBER Working Paper No. 31161. Cambridge, MA: National Bureau of Economic Research.

Burt, R.S. (1987). Social contagion and innovation: cohesion versus structural equivalence, American Journal of Sociology, 92(6), pp. 1287–1335.

Burt, R.S. (1992) Structural holes: the social structure of competition. Cambridge, MA: Harvard University Press.

Davis, F.D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology, MIS Quarterly, 13(3), pp. 319–340.

Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F. and Lakhani, K.R., 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), pp.403-423.

Eloundou, T., Manning, S., Mishkin, P. and Rock, D. (2024) GPTs are GPTs: labor market impact potential of LLMs, Science, 384(6702), pp. 1306–1308.

Gioia, D.A., Corley, K.G. and Hamilton, A.L. (2013). Seeking qualitative rigor in inductive research: notes on the Gioia methodology, Organizational Research Methods, 16(1), pp. 15–31.

Glikson, E. and Woolley, A.W. (2020) Human trust in artificial intelligence: review of empirical research, Academy of Management Annals, 14(2), pp. 627–660.

Goto, M. (2023). Anticipatory innovation of professional services: The case of auditing and artificial intelligence. *Research Policy*, *52*(8), 104828.

Granovetter, M.S. (1973) The strength of weak ties, American Journal of Sociology, 78(6), pp. 1360–1380.

Law Society of Scotland (2024) Navigating AI challenges and risks. Edinburgh: Law Society of Scotland. Available at: https://www.lawscot.org.uk/members/journal-hub/articles/navigating-ai-challenges-and-risks/ (Accessed: 28 July 2026).

Lee, J.D. and See, K.A. (2004). Trust in automation: designing for appropriate reliance, Human Factors, 46(1), pp. 50–80.

Lewicki, R. J., McAllister, D. J., & Bies, R. J. (1998). Trust and distrust: new relationships and realities. *Academy of Management Review*, *23*(3), 438–458.

Li, X., Hess, T. J., & Valacich, J. S. (2008). Why do we trust new technology? A study of initial trust formation with organizational information systems. In *Journal of Strategic Information Systems* (Vol. 17, Issue 1).

Mayer, R.C., Davis, J.H. and Schoorman, F.D. (1995) An integrative model of organizational trust, Academy of Management Review, 20(3), pp. 709–734.

McKnight, D.H., Carter, M., Thatcher, J.B. and Clay, P.F. (2011). Trust in a specific technology: an investigation of its components and measures, ACM Transactions on Management Information Systems, 2(2), Article 12.

McPherson, M., Smith-Lovin, L. and Cook, J.M. (2001). Birds of a feather: homophily in social networks, Annual Review of Sociology, 27, pp. 415–444.

Messerschmidt, C. M., & Hinz, O. (2013). Explaining the adoption of grid computing: An integrated institutional theory and organizational capability approach. *Journal of Strategic Information Systems*, *22*(2), 137–156.

Rogers, E.M. (2003) Diffusion of innovations. 5th edn. New York: Free Press.

Solicitors Regulation Authority (2023) Risk Outlook report: the use of artificial intelligence in the legal market. Birmingham: Solicitors Regulation Authority. Available at: https://www.sra.org.uk/sra/research-publications/artificial-intelligence-legal-market/ (Accessed: 30 July 2026).

Trincado-Munoz, F. J., Cordasco, C., & Vorley, T. (2025). The dark side of AI in professional services. *The Service Industries Journal*, *45*(5–6), 455–474.

Venkatesh, V., Morris, M.G., Davis, G.B. and Davis, F.D. (2003). User acceptance of information technology: toward a unified view, MIS Quarterly, 27(3), pp. 425–478.

6. APPENDIX

Appendix 1 – Coding structure

This section presents the qualitative findings from the case study. Following the Gioia's (2013) methodology, informants' testimonies (first-order concepts) were grouped into researcher-defined themes (second-order themes) and then into overarching aggregate dimensions. The resulting data structure is shown in Figure 1 and provides the backbone for the analysis that was presented in Section 4.2.

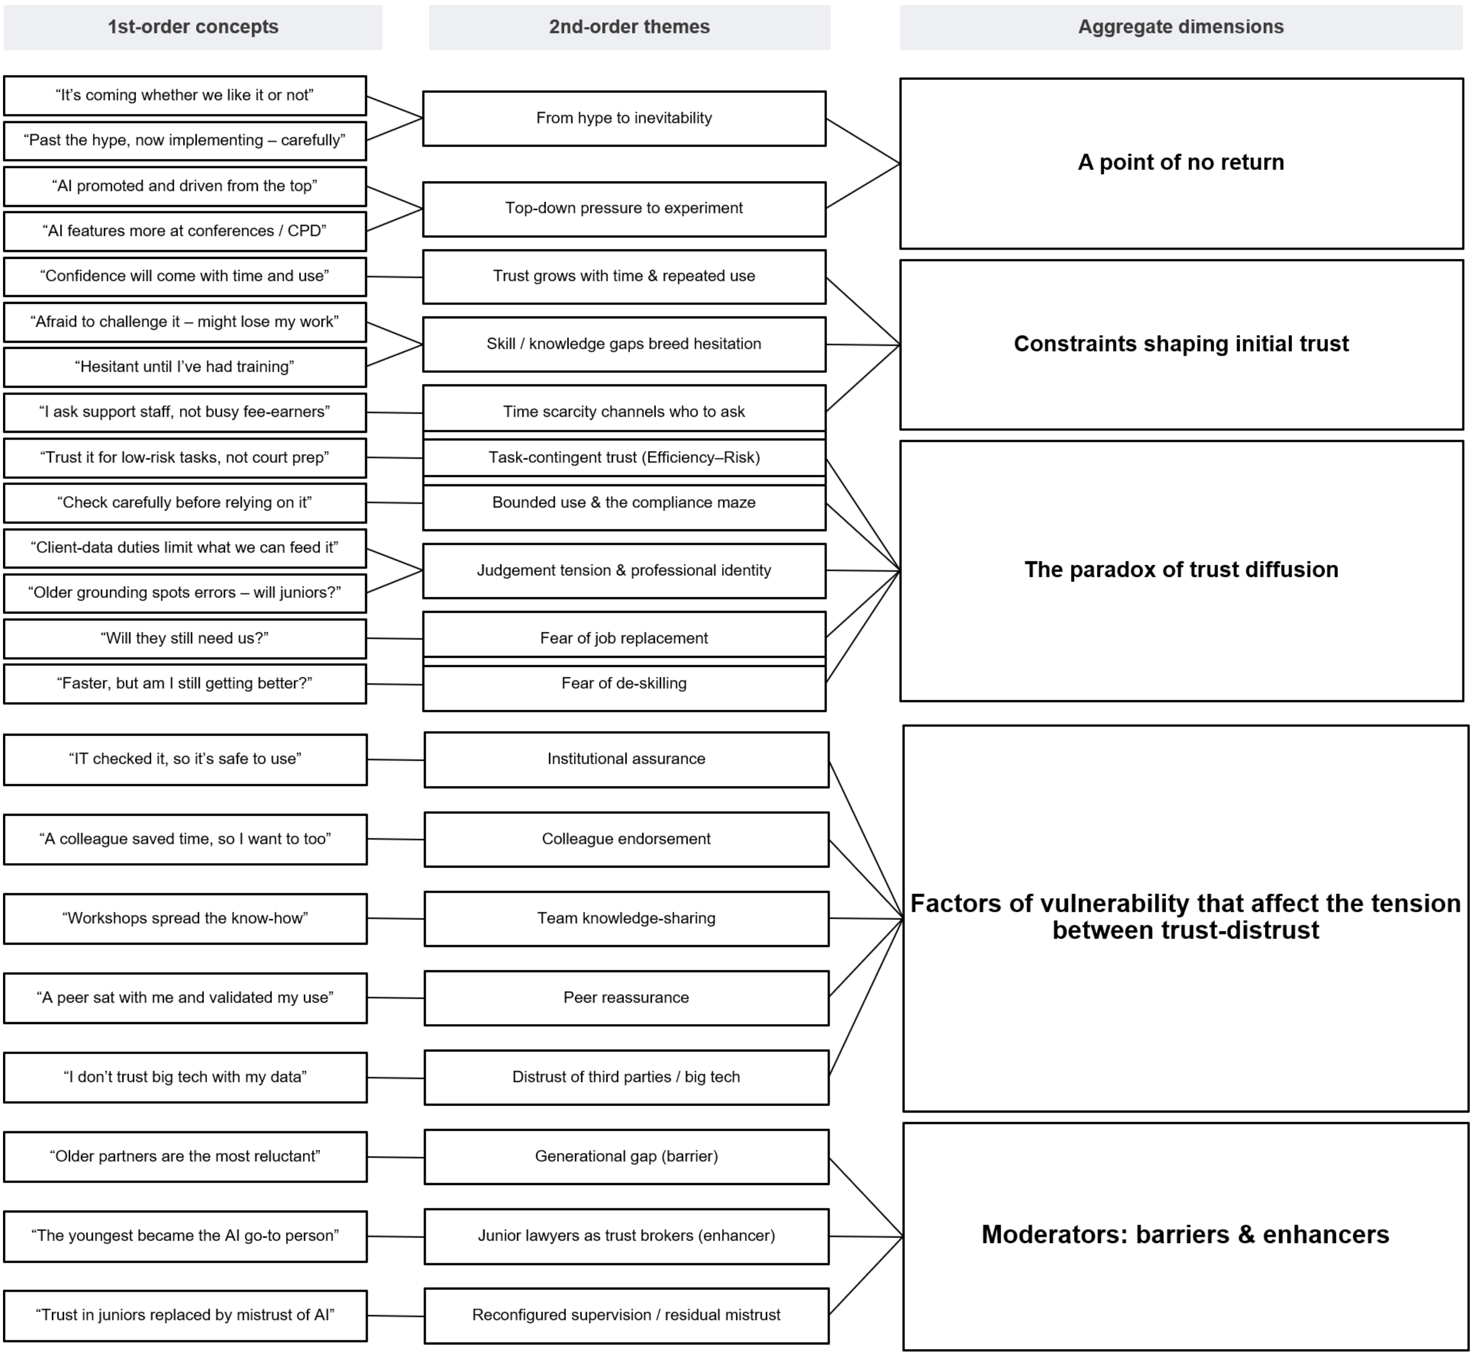

www.ircaucus.ac.uk

Email info@ircaucus.ac.uk Twitter @IRCaucus

INNOVATION & RESEARCH CAUCUS

Delivered with ESRC and Innovate UK