

From AI Adoption to Epistemic Governance: A Sociotechnical View of the Future of Work

Tung Bui¹, Souren Paul², Jim Spohrer³

Abstract

The evolving nature of work is often viewed as a matter of automation, enhancement, productivity increases, or job loss. However, this paper contends that such perspectives, while necessary, are incomplete. As generative, agentic, immersive, and multi-agent AI systems become integral to organizational tasks, they alter the distribution of knowledge, judgment, authority, and accountability. The core question becomes epistemic: who is recognized as the knower, judge, and responsible actor when AI is involved? Drawing on sociotechnical systems theory, human-AI augmentation research, and algorithmic work studies, we introduce epistemic governance—an organization's ability to coordinate AI-assisted value creation with maintaining human judgment, learning, professional identity, and inclusivity. We develop a taxonomy of epistemic relationships between workers and AI, from instrumental support and advisory roles to delegated judgment, collaborative co-agency, authority displacement, and epistemic abstention. We argue that sustainable AI-driven productivity hinges on the organizational governance of AI's epistemic role, rather than on simple adoption. We conclude with a research agenda for IS scholars focused on work design, bypassing developmental constraints, orchestrating multi-agent systems, and fostering inclusive augmentation.

Keywords: Future of work; Augmented intelligence; Epistemic governance; Human-AI collaboration; Agentic AI; Sociotechnical systems.

1. Reframing the future of work: from task change to knowledge authority

Debates about the future of work usually start with questions like: which tasks can be automated, which ones still require human skills, and which jobs are vulnerable? While this perspective helps understand technological changes (Autor et al., 2003), it is becoming less adequate in AI-driven workplaces. Generative AI systems do more than just perform routine tasks; they draft, classify, summarize, simulate, recommend, critique, coach, and increasingly interact with other AI agents. They are now part of work systems not only as tools but also as active participants in interpretation and decision-making.

Work also functions as a knowledge process, where professionals evaluate uncertainty, justify their decisions, and decide what is credible, relevant, and actionable. Epistemology—the study of knowledge and justified belief—gains organizational significance because AI systems now assist in producing claims, explanations, recommendations, and decisions (Fricker, 2007).

This paper builds on the sociotechnical tradition in information systems, which treats technology, organization, work practice, and human agency as mutually shaping rather than separable (Bostrom & Heinen, 1977; Orlikowski, 1992). It also builds on the idea of augmented intelligence: AI should enhance human judgment, learning, and decision-making rather than simply replace human labor, while recognizing that automation and augmentation remain interdependent in practice (Raisch & Krakowski, 2021; Rouse & Spohrer, 2018). We contend that augmentation does not happen spontaneously. It should be a design-and-

¹ Shidler College of Business, University of Hawaii at Manoa, USA

² Drexel University, USA

³ International Society for Service Innovation Professionals (ISSIP), USA

governance choice. Without deliberate governance, the same tools that improve short-term output may erode the human capabilities, identities, and accountabilities that make work meaningful and legitimate.

2. Epistemic relationships in AI-enabled work

We define an epistemic relationship as the arrangement by which humans and AI systems share the act of knowing in work: who generates knowledge, who evaluates it, who is trusted, who may act, and who remains responsible for consequences. Increasingly, AI can create different epistemic relationships across settings. An LLM that corrects grammar is an instrumental tool; the same LLM used to recommend a hiring decision, evaluate a medical risk, or draft a strategic plan becomes part of a more consequential epistemic arrangement. The governance problem, therefore, does not lie in the technology alone. It lies in aligning task, authority, capability, risk, and accountability.

Table 1 proposes a working taxonomy. It is intentionally ordinal rather than deterministic: moving from instrumental to abstention relationships generally increases AI leverage but also raises the risk that human judgment is displaced, de-skilled, or merely symbolic. This logic echoes long-standing concerns about automation misuse, disuse, and skill decay (Bainbridge, 1983; Parasuraman & Riley, 1997), but it is intensified by generative and agentic systems, as AI now produces plausible reasons rather than merely mechanical outputs.

Table 1. A taxonomy of epistemic relationships in AI-enabled work

Type of epistemic relationships	AI role	Human role	Primary governance risk	Illustrative Human-AI interaction
Instrumental	Calculator or tool	Interprets outputs and retains judgment	Low leverage; limited transformation	“AI suggests explanations; I decide what they mean.”
Advisory	Consultant or coach	Evaluates options and chooses	Over-reliance; weak challenge behavior	“I ask AI for alternative strategies, then select one.”
Delegated	Proxy knower for bounded judgments	Monitors exceptions and audits patterns	Accountability gaps; biased delegation	“AI screens reports; humans review edge cases.”
Co-agential	Reasoning collaborator	Co-constructs knowledge with AI	Blurred authorship and responsibility	“I develop a business plan iteratively with AI.”
Authority displacement	Default epistemic authority	Defers unless forced to intervene	Skill atrophy; legitimacy problems	“The AI says this option is the best, so I accept it.”
Epistemic abstention	De facto knower	Withdraws from understanding	Loss of agency and professional meaning	“I no longer understand the process; I just trust what AI says.”

The highest-risk relationships are authority displacement and epistemic abstention. In these cases, the worker no longer uses AI; the worker begins to outsource the very act of knowing. This is both a psychological and organizational risk because accountability, learning, and legitimacy depend on humans' continued ability to understand, contest, and justify AI-enabled outcomes (Lebovitz et al., 2022).

3. The augmentation paradox: productivity with developmental bypass

The first claim is that AI productivity improvements are real but depend on certain conditions. Recent research shows notable gains in specific knowledge-work areas: generative AI support increased productivity in customer service, especially benefiting less experienced workers (Brynjolfsson et al., 2025); experimental studies on professional writing observed faster completion times and higher quality assessments (Noy & Zhang, 2023); and consulting experiments showed significant gains on tasks at the “jagged technological frontier,” but also performance setbacks when users relied on AI for tasks outside of its capabilities

(Dell’Acqua et al., 2023). Therefore, AI’s effectiveness varies by task: it can replace, supplement, or even mislead depending on the work structure and user expertise. Its organizational value depends less on simply adopting AI than on how tasks are structured, the level of expertise, work design, and governance.

The second proposition suggests that poorly designed augmentation could unintentionally bypass essential development. Work is a learning process in which junior analysts improve by drafting, revising, receiving corrections, observing experts, making mistakes, and gradually internalizing tacit judgment (Nonaka & Takeuchi, 1995; Polanyi, 1966). If AI automates too many formative tasks, workers might produce refined results before they have developed the ability to evaluate them. This can speed up organizational processes but weaken the apprenticeship ladder that ensures expertise transfer. This reflects the modern version of the “ironies of automation,” in which the machine takeover of routine tasks leaves humans responsible for exceptional cases, potentially when their practiced competence has declined (Bainbridge, 1983).

This risk is not confined to entry-level jobs. A common assumption is that AI threatens junior workers while senior experts benefit. An alternative possibility is that AI compresses the middle of the occupational hierarchy. If entry-level workers can be augmented by organization-specific AI assistants trained on expert knowledge, some mid-level coordination, synthesis, and drafting roles may become exposed. Conversely, if senior experts use AI to scale their judgment, they may capture greater value. The future of work is thus not about job losses, but a reconfiguration of learning pathways, career ladders, and the distribution of expertise.

4. From individual task performance to epistemic work architecture

The shift from task execution to epistemic orchestration changes the architecture of work. One provocative yet plausible emerging role is the “Student CEO.” The phrase should not be read literally as a claim that every student or novice employee can become an entrepreneur with AI. Rather, it captures a broader organizational possibility: AI enables relatively inexperienced workers to orchestrate a small cognitive workforce. A student or novice professional can use multiple AI “digital workers” to conduct research, synthesize, critique, simulate, design, and present. Faculty and industry mentors, digital twins, and open innovation platforms can further extend this work system.

This model shifts the unit of work from an individual task to a service system: a configuration of people, technologies, data, organizations, and value propositions (Maglio & Spohrer, 2008). The worker is not only a content producer but also an orchestrator of distributed capabilities. The firm may increasingly value not only individual credentials but also demonstrated human-AI micro-enterprises: small teams able to frame problems, mobilize digital agents, compare outputs, prototype solutions, and learn rapidly.

However, the Student CEO metaphor has its limitations in real-world application. Not every valuable worker aspires to lead a team of humans or machines; some excel as deep individual contributors. A humane future of work should avoid prescribing a single heroic archetype of the AI orchestrator. Instead, organizations require diverse work structures: orchestrator roles for those managing agent constellations, expert-maker roles for those advancing their craft, evaluator roles for verification and synthesis, and governance roles to establish accountability across the system. Work design remains crucial because technology influences autonomy, skill utilization, feedback, monitoring, relational quality, well-being, and performance (Parker & Grote, 2022).

5. Epistemic governance as value-creation governance

In this evolving landscape, epistemic governance refers to an organization's ability to appropriately assign judgment, authority, accountability, and learning responsibilities between humans and AI systems. It is a specific aspect of governance related to AI-enabled work. Just as governance gaps occur during digital transformation, epistemic gaps arise when AI systems gain more epistemic influence than organizations can effectively oversee. Over time, this can lead to epistemic governance debt, a hidden vulnerability where

workers increasingly depend on AI outputs without retaining the human capacity to understand, question, or justify those outputs.

The central design challenge is aligning value creation with governance capacity. AI can create value through speed, personalization, coordination, exploration, and service innovation. But value becomes fragile when AI-enabled agency outpaces human oversight, organizational accountability, and worker capability. Research on algorithmic management already shows that algorithms can function as mechanisms of control, evaluation, and discipline rather than neutral tools (Kellogg et al., 2020). AI governance must therefore address not only accuracy and compliance but also the organization of knowledge authority in daily work.

We outline five design principles for operationalizing epistemic governance. These principles translate the abstract concern with knowledge, judgment, and accountability into practical questions for work-system design: clarifying AI's epistemic role, calibrating delegation, preserving apprenticeship, triangulating multi-agent outputs, and ensuring inclusive augmentation (see Table 2).

Table 2. Five design principles for epistemic governance

Design principle	Definition / description	Key questions
Epistemic Role Clarity	Define what role AI plays: tool, adviser, proxy, collaborator, or authority.	What is AI allowed to do? Is its output input, advice, evidence, or decision? Who can challenge it?
Calibrated Delegation	Delegate to AI only when risks, error costs, oversight and accountability are clear.	What can be delegated? What requires human review? Who is accountable when AI fails?
Apprenticeship Preservation	Use AI without bypassing the learning tasks through which workers develop judgment and expertise.	Which tasks remain formative? Can workers explain, compare, and defend AI outputs?
Multi-Agent Triangulation	Govern work involving multiple AI agents by comparing, testing, and reconciling outputs.	Do agents agree for valid reasons? How are conflicts resolved? Who synthesizes the final judgment?
Inclusive Augmentation	Design AI-enabled work so benefits, skills, and opportunities are broadly shared.	Who benefits or loses? Are career paths preserved? Do all workers have access, training, and support?

From our design perspective, ethics should not be viewed as just a compliance step after deployment but as a fundamental performance criterion. Workers are more inclined to critically engage with AI when systems are fair, transparent, open to challenges, and serve their work's purpose. Moreover, systems that seem ethically questionable may lead to resistance, mere passive compliance, superficial use, or misplaced trust. Therefore, in our epistemic view, trust, accountability, motivation, and well-being are essential for sustainable productivity (Glikson & Woolley, 2020; Lee & See, 2004).

6. Research agenda for information systems scholars

We propose four questions for future IS research:

RQ1: How do epistemic relationships vary across occupations, career stages, and organizational contexts? IS research needs finer-grained accounts of when AI is treated as tool, adviser, proxy, collaborator, authority, or substitute for knowing.

RQ2: Under what conditions does AI support capability development rather than developmental bypass? Research should examine how AI alters learning curves, apprenticeship, tacit knowledge, and professional identity.

RQ3: How should organizations govern multi-agent constellations? As workers shift from one AI assistant to multiple agents, research is needed on comparison, auditability, escalation, accountability, and cognitive load.

RQ4: How are the gains from AI-enabled work distributed? IS research should examine not only productivity, but also value distribution, career mobility, inclusion, and who benefits from augmentation.

7. Expected contribution

The expected contribution is threefold. Conceptually, the paper redefines the future-of-work debate from focusing on automation versus augmentation to emphasizing epistemic governance, which includes the design of authority, judgment, accountability, and learning in AI-enabled work. Analytically, it provides a taxonomy of epistemic relationships to compare different human-AI work arrangements. Practically, it proposes design principles aimed at maintaining human capability while harnessing AI-enabled value creation.

The paper seeks to spark discussion around a fundamental question: in an AI-driven economy, what aspects of work must remain uniquely human? The goal isn't for humans to generate every output, but to preserve the ability to judge, learn, contest, govern, and assign meaning. The future of work will not be determined by AI capability alone; it will depend on whether organizations design work systems that extend human capacity or gradually displace the human act of understanding.

References

- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279-1333.
- Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775-779.
- Bostrom, R. P., & Heinen, J. S. (1977). MIS problems and failures: A socio-technical perspective. Part I: The causes. *MIS Quarterly*, 1(3), 17-32.
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889-942.
- Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Kraye, 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. *Harvard Business School Working Paper No. 24-013*.
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627-660.
- 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. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. *Organization Science*, 33(1), 126-148.
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50-80.
- Maglio, P. P., & Spohrer, J. (2008). Fundamentals of service science. *J. of the Academy of Marketing Science*, 36(1), 18-20.
- Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. Oxford University Press.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192.
- Orlikowski, W. J. (1992). The duality of technology: Rethinking the concept of technology in organizations. *Organization Science*, 3(3), 398-427.
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230-253.
- Parker, S. K., & Grote, G. (2022). Automation, algorithms, and beyond: Why work design matters more than ever in a digital world. *Applied Psychology*, 71(4), 1171-1204.
- Polanyi, M. (1966). *The tacit dimension*. Doubleday.
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review*, 46(1), 192-210.
- Rouse, W. B., & Spohrer, J. C. (2018). Automating versus augmenting intelligence. *Journal of Enterprise Transformation*, 8(1-2), 1-21.