

DAMNED IF YOU DO(N'T): HOW EPISTEMIC TECHNOLOGY SPOILS EXPERTISE

Ayman Ajaj

University of Ottawa

aaajaj024@uottawa.ca

Mayur Joshi

University of Ottawa

Mayur.Joshi@telfer.uottawa.ca

Mathieu Bouchard

University of Ottawa

bouchard@telfer.uottawa.ca

ABSTRACT

In both practitioner discourse and academic research, the role of epistemic technologies in shaping the expertise of knowledge workers is well recognized. While the literature has largely focused on the role of epistemic integration in this process, the role of social evaluations remains underexplored. In this study, we focus on a particular type of social evaluation, stigma, and theorize how it reshapes expertise when epistemic technologies are introduced into knowledge work domains. Building on Goffman's foundational work on stigma, we advance a concept of *spoiled expertise* which takes distinct forms based on two dimensions: knowledge workers' stance of adoption of epistemic technology, and locus of evaluation of expertise. Together these two dimensions culminate into four forms of spoiled expertise: (1) *slopped expertise*, (2) *lagging expertise*, (3) *shunned expertise*, and (4) *outdated expertise*. We highlight the main transitions across these forms, and the moral rewriting of expertise under epistemic technologies. Our key contribution is to the literature on knowledge work by advancing a theory of spoiled expertise amid stigma. We conclude by outlining the implications for professionals, and by proposing a research agenda on spoiled expertise in the age of epistemic technologies and algorithmic knowledge production.

1 MOTIVATION AND RESEARCH QUESTION

Epistemic technologies are meant to support the human capacity to produce, acquire, or distribute knowledge, and are increasingly intervening in the very standards by which knowledge claims are made, validated, and acted upon in organizations. Rather than merely codifying established rules or summarizing known indicators, many AI systems generate probabilistic inferences, discover patterns not previously articulated in professional vocabularies, and produce recommendations at scale that can be embedded directly into workflows, dashboards, and governance routines. In doing so, AI does not only assist expert judgment: its output becomes integral to how knowledge claims are produced.

A large body of information systems (IS) and organizational research has examined this tension primarily through a functional-epistemic lens, seeking to explain how technologies reconfigure the division of labor between humans and machines, alter professional jurisdictions and boundaries (Faulconbridge et al., 2024), reshape experts' know-what and know-how (Lebovitz et al., 2021), and generate new forms of knowledge production. Yet another aspect of the link between AI and expertise has been underexplored in literature but is increasingly salient in practice and in public discourse: social evaluations of expertise when knowledge workers adopt or ignore these technologies. In many domains, the controversy is not only whether AI improves decisions, but whether relying on AI signals good or bad expertise.

These evaluative dynamics matter because expertise is not only a repository of skill; it is also a marker of social status that depends on recognition by others. Similar AI technologies can generate divergent patterns of authority and influence across settings, and experts can find themselves in

		Stance of Adoption	
		Adoption	Non-Adoption
Locus of Evaluation	Self	Slopped Expertise	Lagging Expertise
	Others	Shunned Expertise	Outdated Expertise

Table 1: Forms of spoiled expertise

a *damned-if-you-do, damned-if-you-don't* dilemma, penalized whether they adopt AI or resist it. This raises an important theoretical question: **How do social evaluations shape expertise when epistemic technologies are introduced into knowledge work?**

2 STIGMA AS A LENS ON EXPERTISE

Goffman defines stigma as an “attribute that is deeply discrediting”, hence transforming an individual “from a whole and usual person to a tainted, discounted one” and yielding a “spoiled identity”—an actual identity that deviates from the expected social identity (Goffman, 1986, p. 3). A modern extension of this view by Link & Phelan (2001) emphasizes that stigma is not inherent in attributes themselves but arises when audiences collectively interpret certain characteristics as markers of a devalued social identity in a specific context. Within work and occupations, research on “dirty work” (Hughes, 1962) shows how entire professions and their incumbents become stigmatized when the work is perceived as physically or morally tainted, where the attribution of “dirt” is socially constructed rather than inherent in the tasks themselves (Ashforth & Kreiner, 1999).

Stigma also operates in more subtle and interpersonal ways. When an employee carries not-immediately-visible but stigmatizing attributes, they engage in an ongoing management of their social identity (Clair et al., 2005). Stigma-by-association research further demonstrates that negative social evaluations spill over to those who are merely connected to stigmatized targets, although deprived from any stigmatizing attribute (Kulik et al., 2008; Pontikes et al., 2010). We extend these arguments to the context of epistemic technologies and knowledge work. In such context, stigma is not limited to visible marks or low-status (dirty) occupations; expertise itself can become spoiled when stigma attached to a particular technology leads self or others to devalue the scope or trustworthiness of the outputs produced through that expertise.

3 A MODEL OF SPOILED EXPERTISE

The attachment of stigma to knowledge work results in *spoiled expertise*—a form of knowledge whose perceived authority and morality are devalued through processes of impairment. Such expertise has been damaged not by a cognitive deficit or loss of competence, but rather by the spoiling effect of stigma. We conceptualize this spoiled expertise through a two-by-two model (see Table 1). The first axis represents the locus of evaluation, or where the negative evaluation of expertise takes place: expertise appraisal is both self-defined and socially conferred, with knowledge workers appraising their own competence while clients, community members and broader audiences also evaluate whether they are real experts or not. The second axis is the adoption stance, or the position taken by the knowledge worker towards epistemic technologies. Some professionals eagerly adopt and integrate these technologies while others may partially comply, resist, or enact “symbolic conformity by pretending to comply” (Pachidi et al., 2021, p. 19). Crossing these axes yields four forms of spoiled expertise.

Slopped expertise (adoption, self). This form captures experts who adopt epistemic technology and see their own expertise devalued by such adoption. Experts internalize the discourse that epistemic technologies are more objective, accurate, and efficient than humans and, therefore, reinterpret their own judgment and skills as insufficient or intrinsically flawed in comparison. In medical diagnosis, radiologists grappled with AI tools whose opaque outputs diverged from their judgments, hence experiencing “increasing uncertainty” (Lebovitz et al., 2022).

Lagging expertise (non-adoption, self). This form captures experts who do not adopt epistemic technologies, but who internalize societal discourse that valorizes the capabilities and outputs of these technologies. They perceive their non-adoption stance as backward and lazy, even when their actual work performance continues to meet or exceed established organizational standards. Experts may have valid reasons for non-adoption, but they do so under an umbrella of self-doubt and auto-sabotage of their expertise’s worth.

Shunned expertise (adoption, others). This form captures experts who adopt epistemic technologies, and are stigmatized by others for doing so. This adoption is perceived by the expert’s audience as signaling lesser expertise, and the expert as simply an operator or executor whose reliance on technology may indicate a troubling abdication of professional responsibility and weak moral judgment. In law and accounting firms, the potential of machine learning to automate diagnostic work forces professionals into “devaluing machine-learning related tasks” (Faulconbridge et al., 2024, p. 1021) in order to preserve the meaning of the human part of the job.

Outdated expertise (non-adoption, others). This form illustrates experts who do not adopt epistemic technologies and are stigmatized by others for this resistance. These non-adopters become symbols of obstruction to progress and are considered outdated, unwilling to modernize their expertise in line with professional, organizational and societal expectations about what experts should look like nowadays. The coexistence of epistemic technologies with human experts is likely to trigger “clashes over what type of expertise is valued and whose expertise has primacy” (Faraj et al., 2018, p. 65).

4 TRANSITIONS ACROSS THE FOUR FORMS

The four forms of spoiled expertise unfold dynamically over time as technologies stabilize, professional norms shift, and organizational stances about innovation mature. Individuals and occupational groups may move between the four quadrants over time as epistemic technologies diffuse, organizational attitudes and institutional discourse evolve, and jurisdictional boundaries of work change. For instance, early adopters may initially experience internal doubt (*slopped expertise*) but later revalue their hybrid expertise, and shift toward *shunned expertise* if their environment exhibits antagonism toward epistemic technology. Similarly, non-adopters may ignore the obsolescence framing raised by their peers (*outdated expertise*), but as narratives about redundancy gain force they may reluctantly adopt, coming to self-doubt their expertise’s authenticity (*slopped expertise*).

5 THE MORAL REWRITING OF EXPERTISE

The introduction of epistemic technologies into knowledge-intensive settings occasions new interaction “scripts” (Barley, 1986) through which expertise is enacted and evaluated, and expertise becomes spoiled when these scripts are morally coded so that some ways of exercising judgment are treated as appropriate and others stigmatized as inadequate. Stigma also operates through visibility and disclosure, echoing invisible stigma dynamics where individuals decide when and how to disclose devalued identities—as in “techniques of information control” (Goffman, 1986, p. 91). To restore the value of their expertise, experts deploy stigma-management and normalization tactics similar to those documented in dirty work contexts, redefining what counts as real expertise by emphasizing tasks and relational capacities that machines cannot perform—“it’s basically admin, it’s not being a lawyer” (Faulconbridge et al., 2025, p. 2008).

6 THEORETICAL CONTRIBUTIONS AND IMPLICATIONS

Our study contributes to expertise and knowledge-work literature by shifting analytical attention away from performance and efficiency toward the moral and evaluative contexts within which expertise is made visible, credible, or questionable. We also contribute to organization theory by extending Barley’s (1986) notion of technology as an occasion for structuring to explain why identical epistemic technologies diffuse divergently across organizations due to stigma. Finally, our paper contributes to social evaluation research by shifting attention from tainted occupations or individuals to stigma directed at knowledge work guided by technological stances—what one could term *epistemic stigma*.

These dynamics are directly relevant to model accountability in high-stakes settings. In healthcare, clinicians' acceptance or resistance to ML diagnostics is not only technical but morally coded, and accountability regimes can inadvertently spoil expertise: mandates to document overrides may be read as bias or incompetence, discouraging the human oversight that responsible deployment requires. Our model implies concrete design and governance moves—such as incorporating expert reasoning into system design, and avoiding features that publicly tag experts as resistant—so that accountability mechanisms strengthen rather than stigmatize expert judgment.

7 BOUNDARY CONDITIONS AND FUTURE RESEARCH

As with all theory building efforts, our study is subject to clear boundary conditions. It is most applicable to organizational settings in which epistemic technologies are capable of generating and advancing knowledge claims, rather than merely supporting expert judgment: Excel spreadsheets and statistical models are less likely to spoil expertise than a Generative AI system, which allows nearly anyone to perform or claim domain-specific knowledge. Lastly, our study provides several opportunities for empirical work: to trace trajectories across the four forms of spoiled expertise over time; to compare professions such as law, medicine, finance, and HR to reveal whether some fields are more prone to stigmatizing particular technological stances; and to examine how this stigma intersects with existing inequalities linked to gender, race, or organizational status.

REFERENCES

- B. E. Ashforth and G. E. Kreiner. How can you do it? Dirty work and the challenge of constructing a positive identity. *Academy of Management Review*, 24(3):413–434, 1999.
- S. R. Barley. 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, 1986.
- J. A. Clair, J. E. Beatty, and T. L. Maclean. Out of sight but not out of mind: Managing invisible social identities in the workplace. *Academy of Management Review*, 30(1):78–95, 2005.
- S. Faraj, S. Pachidi, and K. Sayegh. Working and organizing in the age of the learning algorithm. *Information and Organization*, 28(1):62–70, 2018.
- J. R. Faulconbridge, A. Sarwar, and M. Spring. Accommodating machine learning algorithms in professional service firms. *Organization Studies*, 45(7):1009–1037, 2024.
- J. R. Faulconbridge, A. Sarwar, and M. Spring. How professionals adapt to artificial intelligence: The role of intertwined boundary work. *Journal of Management Studies*, 62(5):1991–2024, 2025.
- E. Goffman. *Stigma: Notes on the Management of Spoiled Identity*. Touchstone, 1986.
- E. C. Hughes. Good people and dirty work. *Social Problems*, 10(1):3–11, 1962.
- C. T. Kulik, H. T. J. Bainbridge, and C. Cregan. Known by the company we keep: Stigma-by-association effects in the workplace. *Academy of Management Review*, 33(1):216–230, 2008.
- S. Lebovitz, N. Levina, and H. Lifshitz-Assaf. Is AI ground truth really true? The dangers of training and evaluating AI tools based on experts' know-what. *MIS Quarterly*, 45(3):1501–1525, 2021.
- S. Lebovitz, H. Lifshitz-Assaf, and N. Levina. 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, 2022.
- B. G. Link and J. C. Phelan. Conceptualizing stigma. *Annual Review of Sociology*, 27(1):363–385, 2001.
- S. Pachidi, H. Berends, S. Faraj, and M. Huysman. Make way for the algorithms: Symbolic actions and change in a regime of knowing. *Organization Science*, 32(1):18–41, 2021.
- E. Pontikes, G. Negro, and H. Rao. Stained red: A study of stigma by association to blacklisted artists during the “red scare” in Hollywood, 1945 to 1960. *American Sociological Review*, 75(3):456–478, 2010.