

‘Let Me Explain’: A Comparative Field Study on How Experts Enact Authority Over Clients When Facing AI Decisions

Anne-Sophie Mayer^a , Elmira van den Broek^b  and Tomislav Karačić^c

^aLMU Munich; ^bStockholm School of Economics; ^cLondon School of Economics and Political Science

ABSTRACT With organizations increasingly relying on predictive artificial intelligence (AI) technologies for decision-making, experts lose the authority to overrule AI-generated decisions yet remain responsible for presenting them to clients. As experts depend on clients' recognition and approval of decisions, this shift presents a critical disruption to their authority. To investigate how experts respond to this challenge, we adopt a relational perspective that foregrounds the role of audiences in reconfiguring authority. Drawing on a comparative field study, we show how experts sought to reconstruct their authority by engaging in different activities to make clients understand and accept AI decisions, which we call ‘explaining practices’. These practices were shaped by two relational conditions: (1) whether clients recognized the expertise of human experts as unique; and (2) whether interactions between experts and clients provided rich opportunities for learning about clients’ evolving needs. When experts were able to learn about and tailor their explanations to those needs, clients could better make sense of AI decisions and were more willing to accept them, thereby reinforcing expert authority. By contrast, experts who failed to do so left clients with decisions they could not understand or endorse, undermining their authority. This study thereby offers new insights into the complex interplay between expert-client relationships, expert authority, and explaining practices.

Keywords: AI decision-making, artificial intelligence, comparative field study, expert authority, expert-client relationship, explaining practices, relational perspective

Address for reprints: Anne-Sophie Mayer, LMU Munich School of Management, LMU Munich, Ludwigstraße 28, 80539 Munich, Germany (a.s.mayer@lmu.de).

This is an open access article under the terms of the [Creative Commons Attribution](#) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

INTRODUCTION

The use of predictive artificial intelligence (AI) technologies in the workplace presents a fundamental disruption to expert authority by reconfiguring how decisions are made and presented to clients. These technologies rely on machine learning to infer statistical patterns from data to generate predictions and decisions with little – or at times no – involvement from human experts (Berente et al., 2021; Faraj et al., 2018). Moreover, AI outputs are often experienced as opaque by those working with them, as they are typically at odds with human expertise, presenting additional challenges in presenting decisions (Christin, 2020; Lebovitz et al., 2022; Waardenburg et al., 2022). Nevertheless, experts often remain responsible for presenting and communicating such outputs to clients (Klutz and Mulligan, 2019; Li et al., 2024). As such, AI technologies pose distinct challenges around expert authority by questioning the fundamental right of experts to make decisions and their means to present them as understandable and acceptable to clients.

Expert authority refers to the right of experts to issue commands and decisions with the expectation that clients will follow and accept them (Huisng, 2015, 2023; Pakarinen and Huisng, 2025). Attending to matters of authority is critical for management and organizational scholars as it helps explain how professional mandates are carried out and organizational goals are achieved. Experts depend on clients' recognition and acceptance of their decisions for their work to have an impact (Huisng, 2015; Mukherjee and Thomas, 2023). For example, doctors rely on patients adhering to prescribed treatments to fulfill their mandate effectively and contribute to the professional and organizational goals of improving patient health outcomes (Mukherjee and Thomas, 2023). When clients resist, disregard, or question expert decisions, expert authority breaks down and organizational goals are missed (Huisng, 2015). For instance, clients may show undesirable behaviours such as complaining or leaving for other organizations, threatening an organization's goal of satisfying and retaining clients (Huisng, 2015, 2023). As a result, managers may question the value and necessity of expert advice – putting experts' position at stake by, for instance, getting fired, restructured, or degraded (Huisng, 2015). Authority is therefore central to understanding how experts carry out their professional mandates successfully and contribute to wider organizational goals, in which compliance does not occur automatically; it depends on the recognition of clients.

Recent research on AI and work has begun to uncover how experts respond to the challenge of enacting authority over clients despite the reliance on algorithmic and AI technologies for decision-making. These studies highlight that experts often respond to AI disruptions by engaging in practices such as overriding, ignoring, or synthesizing AI predictions with their own expertise to present decisions and advice as competent and legitimate (Anthony, 2021; Christin, 2020; Waardenburg et al., 2022). However, we know little about how experts enact authority when being confronted with AI decisions that cannot easily be overruled, as observed in automated decision-making contexts where algorithms replace experts based on promises of efficiency, objectivity, and fairness (Cecez-Kecmanovic, 2025; Van den Broek et al., 2025). In such cases, traditional means for establishing authority – such as asserting expertise – become less viable, as experts cannot

intervene in the decision itself. Instead, experts must find alternative means to ensure clients accept and follow these decisions. Therefore, our study seeks to explore the following research question: *How do experts enact authority over clients when confronted with AI decisions that need to be presented to clients?*

To explore this research question, we performed a comparative field study on the use of AI in banking, biotechnology, and recruitment where in all cases experts' traditional ways of enacting authority were disrupted by AI decisions. Taking a relational perspective that treats authority as generated, applied, and recognized through interactions with others (e.g., DiBenigno, 2020; Huisng, 2015), we find that experts reconstruct their authority by engaging in so-called 'explaining practices' aimed at making clients understand and accept AI decisions. These practices were shaped by two relational conditions: (1) whether clients recognize experts' expertise as unique; and (2) whether experts' interactions with clients provided rich opportunities for learning about client needs. We find that when experts actively engaged with clients and tailored their practices to clients' needs, clients could better make sense of AI decisions and were more willing to accept them, thereby reinforcing expert authority. In contrast, when experts failed to incorporate client needs into presenting AI decisions, experts left clients with decisions they could not understand or endorse, undermining expert authority.

The relational perspective on expert authority in the age of AI put forward in this study advances our understanding in three ways. First, it reveals how explaining practices provide an alternative pathway to authority when formal roles and expertise do not shield experts from the disruptions posed by AI. Second, it highlights how experts mobilize their expertise differently in these practices depending on client needs and preferences, allowing us to see new opportunities for expertise development. Finally, it demonstrates that both the need for and the ability to contextualize AI explanations are shaped by the relations experts hold with their clients. In what follows, we begin by outlining how a relational perspective differs from dominant perspectives taken on expert authority. We then present the comparative approach followed in this study, paying particular attention to the commonalities and differences between the cases. Following this, we present our separate case analyses and cross-case comparison. We end by discussing the theoretical and practical implications as well as the boundary conditions of our study.

THEORETICAL BACKGROUND

Expert Authority and Expert–Client Relationships

Expert authority is core to an occupation as it solidifies experts' credibility within their domain of work and ensures that clients behave in ways beneficial to occupational and organizational goals (Abbott, 1988; Huisng, 2015). Scholarship on professions and expert occupations has explored how experts develop and maintain authority from different perspectives. One literature stream understands authority as resulting from and closely linked to experts' formal role within a particular institution and respective institutional signals such as training and credentialing (e.g., Greenwood

et al., 2002; Scott, 2008). This line of work highlights how experts are part of certain institutional infrastructures, including professional associations, universities, or government agencies, and how these institutional positions grant them formal authority over a particular domain of work (Carr-Saunders and Wilson, 1933; Larson, 1979). By contrast, a second stream of literature has shifted attention from formal roles and institutional infrastructures toward the role of abstract knowledge in shaping authority (Abbott, 1988; Freidson, 1985; Hughes, 1958). From this view, experts' superior knowledge on how to best achieve outcomes – such as treating a disease or selecting a candidate – grants them the authority to issue commands and decisions to others (Abbott, 1988).

While the above literature streams help identify the role of institutional structures and abstract knowledge in claiming jurisdiction over a domain of work, they provide fewer means to unpack how authority takes shape in relation to specific audiences, such as clients. Attending to the role of clients is crucial given that they are the people experts advise and treat, and it is through clients that experts fulfill their mandate in an organization (Abbott, 1988; Freidson, 1985). Scholars have highlighted that although experts may claim formal authority over specific tasks, their authority over clients is not given but depends on clients' willingness to follow and accept their decisions and advice (DiBenigno, 2020; Huisng, 2015, 2023). For example, Huisng's (2015) study on compliance officers in university laboratories showed that despite their formal position to oversee laboratory compliance of lab researchers, officers had to invest in accommodating, disciplining, and understanding clients to ensure their voluntary compliance.

As a result, an emerging stream of literature has taken a relational perspective that understands expert authority as an interactive process where experts rely on clients' participation and acknowledgment of their decisions as legitimate and competent (DiBenigno, 2020; Emirbayer, 1997; Eyal, 2019; Huisng, 2015). From this view, credentials or abstract knowledge are not sufficient for exercising authority; rather, expert authority depends on the relationships with and responses of those experts treat (DiBenigno, 2020; Eyal, 2019; Huisng, 2015). For experts to secure clients' approval and acceptance of their commands and advice, they must adjust and leverage advice in relation to clients' needs, interests, and values (DiBenigno, 2018; Huisng, 2015; Kaynak and Barley, 2019). For example, public affairs experts, tasked with devising and implementing an organization's political strategy, relied on close interactions with relevant external factors such as community leaders to construct political narratives that persuaded the organization's audience and mobilized support for experts' positions and actions (Kaynak and Barley, 2019).

Scholars taking a relational perspective have shown that various elements of the relationship experts hold with their clients may shape their ability to enact authority over them. For example, when experts interact with clients who fail to recognize the value of their expertise and credentials, they may face additional challenges in issuing decisions and ensuring clients' compliance (DiBenigno, 2020; Eyal, 2019). This was illustrated in the case of mental health experts in the US Army, who were seen as 'out-of-touch outsiders' by the commanders they served, as they lacked knowledge of specific missions and daily work, which often led to their recommendations provoking conflict (DiBenigno, 2020). Similarly, experts serving clients who are strongly

embedded in the organization may face more ‘hostile’ environments to exert authority (Anthony, 2018; Huisng, 2015). For instance, within bureaucratic organizations, experts are often assigned to clients through internal organizational processes, in which clients may not want advice or resent being under the control of experts. These examples illustrate that it is essential to account for the types of relationships experts hold with their clients to understand their ability to gain clients’ approval of their decisions and advice.

AI Technologies and Expert Authority

Recent debates about the use of predictive AI in organizations suggest that these technologies have the potential to transform how experts enact authority by disrupting how decisions and advice are issued to clients (Huang et al., 2019; Lund et al., 2019). Scholars have characterized these technologies by greater autonomy and opacity than prior generations of information systems and argued that, as a result, they pose fundamental changes to expert work (Anthony et al., 2023; Berente et al., 2021; Faraj et al., 2018).

Specifically, AI’s increasing ability to make decisions with little or even without expert involvement challenges traditional foundations of authority by taking over a core task that was previously exclusively performed by experts. While prior work reminds us that it is critical for experts to tailor decisions to clients’ needs to secure their obedience (DiBenigno, 2018; Huisng, 2015, 2023), when AI tools produce such decisions from data, they may radically disrupt the ability of experts to make and adapt decisions in interaction with the clients they serve. For example, loan consultants who had to follow AI decisions experienced less room to negotiate and tailor lending decisions to specific client circumstances by having to follow an automated procedure (e.g., Li et al., 2024; Mayer et al., 2020).

Moreover, the opacity associated with modern AI tools poses additional challenges for experts to enact authority by undermining their expertise in making and explaining decisions to clients. While experts typically emphasize their expertise in relation to clients in order to persuade them of their competence and thereby secure clients’ approval of decisions and advice (Freidson, 1985), AI predictions and decisions are often at odds with expert judgment as they rely on mathematical reasoning and pattern recognition (Burrell, 2016). This makes it difficult or even impossible to understand how and why an AI tool has derived a certain output, making it challenging for experts to interpret, act upon, and explain AI outputs to clients (Anthony, 2021; Waardenburg et al., 2022). For example, Waardenburg et al. (2022) showed how police officers increasingly struggled to make sense of the reports produced by intelligence officers based on opaque AI outputs, leading to growing dependence on those officers to interpret and contextualize the predictions. Consequently, opaque AI predictions can undermine the ability of experts to develop and signify their unique expertise and competencies to clients (Faraj et al., 2018; Pakarinen and Huisng, 2025).

While prior research leads us to expect that experts face critical challenges when their decision-making responsibilities are delegated to AI tools, so far, we lack an in-depth understanding of how experts enact authority over clients despite these disruptions. Those studies that have considered clients have largely focused on situations where experts still

possessed a large degree of discretion to consider or override the AI prediction (e.g., Anthony, 2021; Waardenburg et al., 2022). However, as experts are increasingly discounted and taken out of the loop due to their perceived cognitive limitations and potential biases (Cecez-Kecmanovic, 2025; Klutz and Mulligan, 2019; Strich et al., 2021; Van den Broek et al., 2025), experts find themselves facing AI-generated decisions that cannot be dismissed but still need to be presented and justified to clients. For example, lawyers need to explain their strategy to clients and justify decisions toward judges despite AI tools automating and assisting lawyerly decision-making (Klutz and Mulligan, 2019). Drawing on a relational perspective (Anteby et al., 2016; Emirbayer, 1997; Pakarinen and Huisng, 2025), we set out to examine how experts enact authority when confronted with AI decisions that cannot be overruled but need to be explained to clients.

METHOD

This research builds on insights from a comparative field study in three settings where higher management introduced an AI tool to make decisions previously performed by human experts. Comparative field studies provide deeper insights into the underlying mechanisms, processes, and contexts that shape phenomena of interest (Anthony et al., 2023; Bechky and O'Mahony, 2015), and have proven fruitful in studying novel phenomena in management and organizations (Bechky and Okhuysen, 2011; Levina and Vaast, 2005; Staudenmayer et al., 2002).

We came to select our settings – banking, biotechnology, and recruitment – based on regular case discussions among the authors. All authors were part of the same research group while working on separate field studies on the use of AI in expert work. In the course of these conversations, we stumbled upon a compelling commonality between our three cases, which neither of us was studying separately: experts' confrontation with AI tools that generate decisions for clients that cannot easily be overruled but still have to be explained. At the same time, the cases differed in how this disruption played out for experts, offering a valuable opportunity to explore theoretical differences. These similarities and differences motivated our comparative study into how experts seek to enact authority over clients despite the challenges posed by AI. Table I presents an overview of the different cases.

Data Collection

Data were collected from separate field studies, with each author collecting in-depth data on one of the cases. Although data collection was performed separately, all cases shared the aim of developing an in-depth understanding of how the use of AI in decision-making changed the work experts performed in relation to their clients. Thus, this study primarily focuses on experts' interactions with clients around AI-generated decisions. All researchers were guided by ethnographic methods, which involved field observations, semi-structured interviews, and analysis of archival data. Table II gives an overview of our data sources and their use in analysis.

Banking. The first author collected data at a large German bank with about 135,000 employees between January 2019 and January 2025. The data included field observations

Table I. Case overview

	<i>Banking</i>	<i>Biotechnology</i>	<i>Recruitment</i>
Professional group	Loan consultants	Seed sorters	Recruiters
Key clients	Customers	Supply chain managers	Senior managers
Decision made by AI	Loan approval	Seed selection	Candidate selection
AI outputs supporting the AI decision	One-pager with AI decision rules that approximate why a loan is rejected	Graphs and feature importance plots that approximate why a seed is selected	Personality scores and graphs with personality traits that approximate why a candidate is selected

that focused on the daily work of loan consultants as well as respective interactions with their core clients, loan customers. Moreover, the researcher engaged in 70 semi-structured interviews and informal conversations with loan consultants and customers. As a complement, data collection involved the review of company documents, such as brochures and presentations.

Biotechnology. The third author conducted fieldwork at a biotechnology company, a global market leader in developing new plant varieties with about 3000 employees, from March 2021 to April 2023. The researcher conducted field observations, shadowing seed sorters over full working days and documenting their daily work practices and interactions with supply chain managers, their core clients. In addition, the researcher conducted 23 semi-structured interviews with seed sorters and supply chain managers. Finally, company documents were collected, such as internal reports, brochures, and presentations.

Recruitment. The second author gathered data at the HR department of a global fast-moving consumer goods firm with over 200,000 employees from October 2018 to April 2022. The researcher conducted field observations, shadowing recruiters over full working days and attending meetings with their core clients – senior managers – particularly job interviews where recruiters presented chosen candidates. The researcher also engaged in 46 semi-structured interviews and informal conversations with recruiters and senior managers. As a complement, company documents, such as PowerPoints, AI outputs, and meeting notes were collected.

Research Settings

We now briefly explore the commonalities and differences between our cases, thereby laying the foundation for understanding how experts enacted authority over clients when confronted with AI decisions (Table III).

Two commonalities in how expert authority became disrupted make these cases relevant and unique. First, in all cases, experts lost decision-making responsibilities to AI,

Table II. Cross-study comparison of methods

Data types	Cases	Data range	Use in analysis
Field observations	Banking	~31 days	Provided insights into how loan consultancies were performed with AI and how loan consultants explained loan decisions to customers.
	Biotechnology	~170 days	Provided insights into how seed sorting was performed with AI and how seed sorters explained seed selections to supply chain managers.
	Recruitment	~180 days	Provided insights into how candidate selection was performed with AI and how recruiters explained chosen candidates to senior managers.
Semi-structured interviews and informal conversations	Banking	55 loan consultants 15 customers	Provided insights into: <ul style="list-style-type: none"> loan consultants' work processes and practices before and with AI. loan consultants' challenges in explaining loan decisions to customers with AI. customers' reactions to provided decisions and explanations.
	Biotechnology	15 seed sorters 8 supply chain managers	Provided additional insight into: <ul style="list-style-type: none"> seed sorters' evaluation and selection practices before and with AI. seed sorters' struggles in explaining seed selections and their ways of coping with the emerging issues. supply chain managers' critiques and needs concerning seed selection and subsequent explanations.
	Recruitment	35 recruiters 21 senior managers	Provided broader insight into: <ul style="list-style-type: none"> recruiters' assessment and selection practices before and with AI. recruiters' challenges and strategies in explaining chosen candidates to senior managers. senior managers' issues and experiences around chosen candidates and subsequent explanations.
Company documents	Banking Biotechnology Recruitment	AI outputs and predictions, brochures, presentations	Provided insights into the background and outputs of the AI tool as well as intended goals.

with experts facing decisions generated by AI that could not easily be overruled. Experts were still charged with overseeing and presenting these decisions to their clients. Second, all experts experienced these AI-generated decisions as largely opaque, given that they were grounded in machine learning models that were at odds with their own expert judgment. To support experts in presenting AI decisions to clients, all AI tools generated supportive outputs in the form of written reasons, graphs, or features underlying the decision. Appendix A presents an overview of the outputs experts received to explain AI decisions.

Two theoretical differences in the relationship between experts and clients were important for understanding how experts responded to the AI disruption of their authority. First, our cases differed in terms of whether clients recognized their expertise as unique. In banking and biotechnology, clients understood experts' knowledge as distinct and valuable for making decisions, whereas in recruitment, clients shared related expertise in the domain, viewing recruiters' judgment as less valuable. Second, the cases diverged in terms of whether there were opportunities for experts to learn from and respond to client needs through rich and continuous interactions. In banking, experts typically had one-off exchanges with clients, with limited opportunities for learning and feedback. In contrast, experts in biotechnology and recruitment continuously and closely engaged with the same group of clients, exposing them to rich learning opportunities.

Data Analysis

Inspired by comparative ethnographic studies that bring rich insights into underlying patterns, themes, and variations (Bechky and Okhuysen, 2011; Levina and Vaast, 2005; O'Mahony and Bechky, 2008), we compared and analysed the data from the three selected settings to explore the dynamics between AI decisions, expert-client relations, and expert authority. This comparative methodology differs from multiple case studies (Eisenhardt, 1989; Yin, 1992) in that there are no a priori categories used in collecting data to frame the initial analysis. Instead, we studied similarities and differences among the categories independently developed in each study (Bechky and Okhuysen, 2011).

Our data analysis was initially inspired by constructivist grounded theory methods (Charmaz, 2006), which involved going back and forth between our collected data and theoretically emerging concepts. Each author began with an emic analysis of the data, independently coding the practices experts engaged in to present AI decisions to their clients. When we compared the initial insights, we noticed interesting differences between cases. For example, in banking and biotechnology, we discovered that experts initially refrained from showing any outputs produced by the AI tool when communicating decisions to their clients. In recruitment, recruiters prominently featured AI graphs and scores in their explanations to clients. We found these differences to be an intriguing empirical puzzle that pushed us toward unpacking how and why experts enacted different practices in presenting AI decisions.

We proceeded with a systematic analysis of similarities and differences in how experts coped with the need to communicate AI decisions to their clients. During this step, we met

Table III. Commonalities and differences between the three cases

Aspect	Banking	Biotechnology	Recruitment
Commonalities			
Decision-making autonomy of AI	<i>AI tool makes decisions instead of experts:</i> Loan decisions are made by AI instead of loan consultants	<i>AI tool makes decisions instead of experts:</i> Seed selections are made by AI instead of seed sorters	<i>AI tool makes decisions instead of experts:</i> Candidate selections are made by AI instead of recruiters
Opaqueness of AI decisions	<i>Presenting opaque AI decisions to clients:</i> AI-generated loan decisions that are at odds with consultants' expertise need to be presented to clients	<i>Presenting opaque AI decisions to clients:</i> AI-generated seed selections that are at odds with seed sorters' expertise need to be presented to clients	<i>Presenting opaque AI decisions to clients:</i> AI-generated candidate selections that are at odds with recruiters' expertise need to be presented to clients
Differences			
Recognition of expertise as unique by clients	<i>Recognized expertise:</i> Loan consultants benefit from expertise that is recognized as unique by clients	<i>Recognized expertise:</i> Seed sorters benefit from expertise that is recognized as unique by clients	<i>Questioned expertise:</i> Recruiters share complementary expertise with clients that is questioned as unique
Learning about client needs through client interactions	<i>Lack of learning through sparse interactions:</i> Loan consultants typically have one-off exchanges with clients	<i>Learning through rich interactions:</i> Seed sorters have continuous and deep exchanges with clients	<i>Learning through rich interactions:</i> Recruiters have continuous and deep exchanges with clients

frequently and discussed narrative descriptions, codes, and emergent categories in an effort to identify what activities experts engaged in to help clients understand and accept AI decisions, which we labelled 'explaining practices'. We categorized these practices based on how experts engaged with the outputs produced by the AI tool and the role their expertise played in explaining decisions. In particular, we noticed that some activities centred on hiding the AI outputs and foregrounding expertise in explaining decisions, while others prominently featured the outputs. We referred to these practices as *masking* and *showcasing*, respectively. Additionally, we identified activities focused on combining AI outputs with expertise – a practice we termed *enhancing*. In contrast, other activities involved tweaking, refining, and fine-tuning the AI outputs through technical means, which we labelled *calibrating*.

As we realized that expert groups engaged in different explaining practices, we started to explore how and why these practices differed. Inspired by the literature on professions and expert occupations (e.g., DiBenigno, 2020; Huisng, 2015; Prasad, 1993), we found that specific differences in the relations experts held with their clients were consequential in shaping these practices. Those experts who held expertise typically recognized as unique and

valuable by their clients engaged in practices that largely obscured the role of AI while still emphasizing their expertise, albeit not consequential anymore in decision-making. By contrast, experts who faced clients who largely questioned their expertise as unique and valuable strongly embraced the AI tool and its outputs in explaining decisions to clients.

As we began to notice shifts in how experts explained AI decisions to clients, we turned to temporal bracketing to analyse these patterns over time (Langley, 1999). We identified discontinuities in expert-client interactions that signalled a change in how experts approached the task of explaining. For instance, in the seed sorting case, we observed how experts initially deflected expert questions by using vague references to seed colour, while gradually shifting toward actively using AI outputs in client conversations. Such shifts marked critical turning points in practices that we used to delineate temporal brackets. Zooming in on the activities, issues, and interactions preceding and following these turning points enabled us to see how specific explaining practices in one period created tensions or mismatches with clients' needs that experts responded to in the next. Importantly, we found that experts' ability to recognize and respond to these tensions depended on opportunities for interacting with and learning about client needs. Those who benefited from rich feedback – through direct questions, pushback, or requests for clarification – were better able to adjust their explaining practices to gain clients' approval. In contrast, others remained locked in patterns that failed to meet client needs and ultimately undermined expert authority. In the following sections, we detail these evolving trajectories and their consequences across cases.

FINDINGS

Below, we present each case separately, illustrating how experts initially sought to enact authority over clients, how the introduction of AI disrupted this authority and prompted specific explaining practices, and how these practices evolved over time through their interactions with clients. We conclude the separate case descriptions with a cross-case analysis.

Banking Case

Traditional situation: Experts enact authority over clients by relying on their expertise. In banking, we observed how loan consultants traditionally enacted authority over clients by relying on their expertise, which helped persuade customers to view loan decisions as reasonable and acceptable.

Traditionally, loan consultants enjoyed significant discretion in issuing loan decisions and respective conditions. In particular, loan consultants were able to tailor loan decisions and conditions to customers' individual situations, for example, when agreeing to grant a loan in borderline cases but at a higher interest rate. While operating under standardized industry criteria (e.g., no loan was given to unemployed or over-indebted customers), consultants heavily relied on their expertise and own judgment when deciding on loan approvals and tailoring them to client needs. For example, when evaluating the creditworthiness of a customer, consultants took into consideration their experiences with past customers that had similar backgrounds regarding expenses, income, and life circumstances:

Over the years, you get a feeling for what works and what doesn't. For example, if you see that the monthly expenses vary a lot, you get suspicious. Especially if the account limit is sometimes overdrawn, this is usually an indicator that the customer does not have a good relationship with money. (Loan consultant)

Loan consultants took great care in explaining loan decisions to clients by tailoring their communication to customers' individual situations and preferences. Their explanations ranged from verbal justifications to using visual aids such as charts or calculations, depending on 'each customer's background and need for additional information or explanations' (Loan consultant). In explaining loan decisions, consultants often pointed out clear and concrete reasons behind the decision—particularly financial indicators that customers could easily grasp, such as high levels of existing debt. Doing so allowed them to present decisions as reasonable and acceptable, even in the case of loan rejections, by appealing to well-recognized expert standards and the customers' own financial well-being. For example, one consultant reflected on how she explained a rejection as reasonable to customers by framing it around the risk of over-indebtedness:

I've always said that we are on the side of the customer, so if possible, we want the loan [approval] for the customer, we want to help. But I've also clearly said: 'Look, if you can't afford the loan, it's for the best to say no. Otherwise, you'll end up in a downward spiral of debt.' And I think customers value that because they know we are on their side. (Loan consultant)

By grounding their explanations in tangible financial reasons and expert terms, consultants helped ensure that even unfavourable outcomes could be interpreted as in the customer's interest, in which customers generally submitted to their advice and decisions.

Initial explaining practice: Masking AI outputs as a result of interpreting AI as a threat. Once a predictive AI tool was introduced for issuing loan decisions, it presented an important threat in the eyes of consultants to their expertise, which was traditionally core to how they enacted authority over customers. While previously, loan consultants made and presented tailored loan decisions by relying on their expertise, the AI tool now issued these decisions based on data-driven patterns that were often at odds with their judgment. In an attempt to maintain their authority despite the disruption of AI, consultants masked the AI outputs in explaining loan decisions to customers.

In 2017, the bank's management decided to fully rely on an AI tool that promised to autonomously predict customers' creditworthiness by detecting patterns in vast datasets, including historical customer behaviour and data. Management instructed loan consultants to merely enter customers' data into the AI tool and click on 'make a decision'. Machine learning models then computed and displayed whether the loan was granted (and on which conditions) or rejected. In case of a loan rejection, the AI tool produced a one-pager with bullet points summarizing the underlying reasons to help consultants understand and explain the decisions to their clients. Higher management reasoned that

this AI-driven decision-making process would be more consistent and objective than their traditional process led by consultants, thereby contributing to lower default rates.

With the introduction of AI, loan consultants' decision-making responsibilities fundamentally changed as they lost their discretion in deciding whether to grant a customer a loan and under what conditions (e.g., interest rates). Critically, the loan decisions as well as the one-pager produced by the technology were often not intuitive or comprehensible to consultants. For instance, a loan consultant was puzzled when the system rejected an application that, based on his expert judgment, should have been approved:

Sometimes these decisions don't make any sense. I mean how is it possible that someone with an above average income and no debts doesn't get a loan? This just doesn't make sense. And then the [one-pager] says: unstable financial situation. What is that supposed to mean? (Loan consultant)

At the same time, loan consultants remained responsible for explaining the AI decisions to their customers. This placed them in the difficult situation of explaining loan decisions to customers that were no longer made by them or grounded in their expertise.

Experts interpreted the AI disruption as a threat to their authority over clients, as customers traditionally valued loan consultants' banking expertise and perceived them as 'the expert' (Bank customer). Given that the loan decisions were no longer grounded in their expertise, loan consultants feared that their expertise would become obsolete in the eyes of customers as their role became reduced to entering customers' data into the AI tool:

I didn't learn this job just to sit there in a purely accompanying role. Because well, the danger is really that the customer takes over these simple tasks by themselves and that we don't really have anything to do anymore. (...) Yes, for the employees, there is the fear of losing their job in the long run. Definitely. So, I think that the employees really have a problem communicating this [use of and dependence on AI] to (...) the customers because I am, in a way, rationalizing myself out of the equation. I see this as the biggest issue for employees. (Loan consultant)

In loan consultants' attempt to uphold their authority and thus, to remain valuable in the eyes of customers, consultants continued to make their expertise central in presenting AI decisions to customers. This practice, which we came to label as *masking*, involves experts foregrounding their expertise in explaining AI decisions while hiding AI outputs in front of customers.

Specifically, consultants decided to withhold the bullet-point summary provided by the AI tool, for instance, by not turning the computer screen to the customer, and explaining the reasons for a loan rejection or approval based on their own terminology and reasoning. In particular, in cases where loan consultants were puzzled by the AI decision, they tried to come up with alternative explanations based on their own expertise. For example, if the one-pager stated a too-low income as a reason for rejection, but consultants reasoned that the customer had an above-average income, they came up with their own generic explanations grounded in their previous experience

of common causes for loan rejections, such as a new income threshold or a negative credit history:

It's quite convenient to say that because of inflation or changes in the currency market, there are new thresholds in income, assets, whatever. How should the customer know? You can always argue that this threshold is there to protect the customer from over-indebtedness and to ensure that he is really able to repay the loan. (Loan consultant)

Experts engaged in this practice as they felt that explaining in competent ways why these decisions were made was a key way to demonstrate their relevance to customers:

We increasingly become dependent on technology in our job, but you don't want to admit that to the customer, right? And you need to highlight that there is still a value in seeing your personal consultant rather than doing it [applying for a loan] online. (Loan consultant)

Demonstrating experts' continued relevance to customers was crucial for loan consultants because, as the AI tool took over core tasks that customers could potentially perform on their own, consultants risked being perceived as redundant and ultimately 'rationalizing [themselves] out of the equation'. Thus, consultants were strongly committed to emphasizing that their expertise was still valuable for customers.

Stagnating explaining practice: Continuing to mask AI outputs due to a lack of learning from clients. Although these explanations allowed consultants to use and emphasize their expertise toward customers, customers often experienced these explanations as overly generic and at odds with their individual financial situations. For example, during one observed instance, the consultant looked at the computer screen and reported that the customer's loan application was, unfortunately, rejected due to a problematic credit history. The customer looked surprised and said: 'That's strange, I don't have any outstanding debts or so, so no idea what should be in [the credit history]' (Bank customer). Still, the consultant kept arguing that the customer must have a negative entry in his credit history that he might be unaware of. In response, the customer looked puzzled, paused, and then said 'Mhm, okay, weird' (Bank customer). Hence, customers often struggled to follow the explanations presented by consultants and were not genuinely convinced by the final decisions.

Despite the confusion and struggles experienced by customers, consultants had few opportunities to learn about customers' needs and improve their explanations. Consultants typically only interact with loan customers during one-off exchanges, as the same customer typically applies 'only once or a few times over several years for a loan' (field note). During these sparse interactions, customers often felt that they, as laypeople, lacked the appropriate knowledge to challenge consultants who benefit from unique domain expertise. As a result, they refrained from openly expressing their dissatisfaction or concern with the provided explanation. For instance, one customer expressed: 'Okay, if you say so; I guess you know best' (Bank customer). By not being contractually bound to one

bank, customers perceived it as most convenient to simply seek advice from a different bank when experiencing loan decisions as implausible and disconnected from their own perceived financial situation: 'Maybe I just don't get it, but maybe I'll try it at [competitor bank]' (Bank customer).

As a result, loan consultants largely continued their masking practice when explaining AI decisions to customers, as they did not perceive the need to adapt their practice due to a lack of critical feedback from clients. However, as observed during a final field visit in 2025, higher management reflected on this development as increasingly problematic because it became more difficult for the bank to continue demonstrating the benefits of personal loan consultancies – a core element of the bank's business model – when consultants offered limited value in making and explaining decisions:

It's not easy to leverage the benefits of automated technologies like [AI tool] for us and customers but not to run into a situation where customers feel everything is done by the technology and it's easier to directly go to a digital bank because there is no value anymore in a bank that relies on physical branches and personal consultancies. (Bank manager)

The manager's concern revealed how masking may backfire – rather than reassuring clients, it may highlight consultants' diminishing relevance, ultimately undermining the consultants' authority and the bank's traditional value proposition.

Biotechnology Case

Traditional situation: Experts enact authority over clients by relying on their expertise. In biotechnology, we observed how seed sorters relied on their embodied expertise to establish authority over supply chain managers, their key clients. While supply chain managers at times interrogated seed sorters' judgments by asking clarifying questions, they generally recognized and relied on the expertise seed sorters developed through prolonged and direct physical experience with seeds.

Traditionally, the role of seed sorters involved making decisions concerning the quality of seeds and how to improve them, after which supply chain managers organized the packaging and shipping of seeds based on the seed sorters' sorting advice. Specifically, seed sorters inspected the seeds based on 'decades of working with the seed, using [their] eyes and hands' (Seed sorter). They would place their hands into bags of seeds, rub them in their hands, and look at the seeds in their palms to separate the low-quality from the high-quality seeds. After inspection, seed sorters determined which of the 150 sorting parameters (e.g., colour and shape) the seed would be sorted using a variety of machines. Before sorting, they estimated the number of seeds that would be sorted and the quality of the selected seeds, which would be reported to their clients, the supply chain managers. As a seed sorter explained, their expertise was highly valuable for the organization:

Sorting is hard work and you need to make a lot of decisions, which in the end mean we either earn money or we lose money and it all comes down to knowing the seed. (Seed sorter)

Supply chain managers typically agreed with the experts' sorting strategies, trusting the seed sorters' deep, hands-on expertise and believing seed sorters 'know what they are doing when it comes to seed' (Supply chain manager). One reason for this compliance was that seed sorters provided insights based on their expertise in evaluating overall quality and specific seed defects, which helped managers determine for themselves which seeds should be sent to which buyer. Occasionally, managers requested additional explanations and justifications for specific decisions, particularly when expert advice complicated seed distribution. For example, one seed sorter described how a supply chain manager had emailed him to question the removal of a large portion of a tomato batch. The manager was worried there would not be enough seed left to fill an Eastern European order and asked if the defects really mattered. While expert decisions were sometimes questioned and required ongoing work by seed sorters, supply chain managers relied on them. As one manager explained:

Seed sorters really need to tell us some information that helps us organize the logistics, like when we decide who gets which seed, so they tell us what is actually the problem with the seed. (Supply chain manager)

By relying on their own eyes and hands to make such assessments, seed sorters communicated these nuances to supply chain managers, reinforcing the value of their embodied expertise. As a result, in the seed sorting case, experts leveraged their unique, embodied expertise gained through prolonged and direct physical engagement with seeds to establish authority over supply chain managers.

Initial explaining practice: Masking AI outputs as a result of interpreting AI as a threat to expertise. The introduction of AI presented a significant change for seed sorting practices, disrupting the embodied expertise through which seed sorters traditionally enacted authority over their clients. In particular, we observed how seed sorters perceived losing decision-making power to the technology, thereby threatening their authority over clients. Similar to the banking case, experts initially aimed to protect their authority by *masking* AI outputs while foregrounding their expertise in explaining seed selections to clients.

In 2017, the company board decided to invest in the development of an AI tool to automate the sorting of high-quality seeds, which traditionally involved the core decision-making responsibility of seed sorters. The AI tool combined x-ray, chlorophyll, and light-based imaging with machine learning models trained to distinguish high- from low-quality seeds. As an AI developer explained:

The idea was to automatically know for each individual seed whether it is good or not. Instead of guessing [by seed sorters], the machine actually 'sees' [the quality of a seed]. (AI developer)

The introduction of the AI sorting machines led to a radical shift in seed sorters' decision-making responsibilities. Seed sorters were no longer involved in evaluating and making decisions regarding the sorting of seeds. Instead, they acted as operators,

responsible for turning on and off the sorting machines, as well as for the manual work of loading the seeds into the machine and cleaning the work floor. As the machines sorted the seeds immediately after evaluating them, there was no opportunity for experts to overturn the decisions. To provide insights into how the seeds were evaluated and sorted, the AI tool produced a report that presented the computed dominant features of a particular batch and a sorting graph (see Appendix A). Given that seed sorters oversaw the sorting machines, management's intention was that they also provide explanations of sorting outcomes to supply chain managers, who lacked details on how the tool worked based on the sorting graph.

Yet many seed sorters struggled to reconcile the AI outputs with their embodied understanding of seed quality. Watching the machines in operation, they often found it difficult to interpret why specific decisions had been made. One sorter remarked:

Sometimes, I just sit behind the machine, watch the images and seeds go by, and see if I still agree with the report and the seed being sorted out. Sometimes, I don't agree, but then I hear from buyers that it is good, meaning the model does it much better than I do, and I don't know how it is doing it. (Seed sorter)

This growing sense of disconnection from the decision process led seed sorters to fear that their expertise was being rendered obsolete. While seed sorters 'knew what was going in the bag and why' at the time when they relied on their own bodies to make expert judgments, now they 'would just give them advice based on [familiar categories]', aiming to say 'these are the results of our sorting, and here's what it means for you.' (Seed sorter) When asked about this practice, seed sorters often expressed concern over the lack of understanding of how the model works. As one seed sorter noted:

'[I was never sure] how did this [model] decide that? And then if someone from supply chain asks, you feel a bit stupid. (Seed sorter)

In response to the disruption that the AI tool posed, seed sorters, like loan consultants, engaged in the explaining practice of *masking*. Rather than openly presenting the AI-generated reports, they drew on familiar embodied language to craft plausible explanations for the observed outcomes. Their goal was to preserve the impression that their expertise remained central to seed evaluation, even as the AI models took over decision-making.

In practice, this involved presenting AI outputs in terms that resonated with their prior experience. For example, when a batch was classified as low quality, sorters would visually inspect the seeds and attribute the result to visible defects they identified, without knowing whether the AI model made the decision based on the same criteria. As one seed sorter explained:

You can't just tell them the model made it this way. You say something about the color or the size, even if that was not the reason from the model. Otherwise they don't know what to do with it. (Seed sorter)

These explanations appeared initially sufficient for supply chain managers ‘as long as it makes sense for what [supply chain managers] need to do next [with organizing logistics]’ (supply chain manager). Such practices allowed seed sorters to maintain a sense of relevance and to signal that their embodied expertise still informed sorting outcomes.

Shifting explaining practice: Enhancing AI outputs by learning from clients. Although clients respected these initial explanations, seed sorters were faced with a challenge of their explanations lacking the contextual specificity that supply chain managers required. In the seed sorting case, the limitations of masking practices became apparent through repeated, task-specific interactions with supply chain managers. When seed sorters offered generic explanations, often citing broad reasons like ‘colour’, which they could identify using the traditional way of evaluating seeds, tensions began to surface. In one case, a manager reviewing a batch report called a sorter to ask why so many seeds had been rejected. The sorter hesitated, glanced at the AI graph, and replied that it was probably due to discolouration. Later that week, the manager complained in a team meeting that ‘sometimes they just tell us the batch is not good because of color, but we need more than that to plan where to send the seed’ (supply chain manager). This was evidence of a broader issue raised by supply chain managers:

They cannot always tell us what went wrong with the seed. They need to give specific points about the seed so we can see to which country it would be best to send this batch. (Supply chain manager)

The ongoing conversations between seed sorters and supply chain managers made it increasingly clear when explanations were not sufficient, but the shift in practice did not come from pressure alone. Sorters were already working closely with the AI system in their daily routines. They monitored outputs, loaded batches, and observed how different seed lots were processed. Over time, they began to notice recurring features in the images and graphs that hinted at subtle differences in seed quality, such as slight changes in colour, texture, or internal structure. These features were not part of their earlier sorting criteria but became meaningful through repeated observation and informal conversations among colleagues. One sorter explained:

I kept noticing small things and started wondering if that meant slower emergence or uneven plants. There were no answers. We had to study the seed and the machines ourselves and find it out. (Seed sorter)

As their interpretations developed, sorters began to use these insights when speaking with supply chain managers. Instead of offering general explanations, they referred to specific parts of the images and related them to growth performance or client needs. One seed sorter noted:

Before I would just say it’s color, but now I can show them. This part here might mean some seeds come up late. (Seed sorter)

The shift led to a new explaining practice of *enhancing*: developing novel insights beyond what the AI tool offered to make decisions more explainable and acceptable to clients. The move from masking to enhancing emerged from the sorters' interaction with clients, where they shared their own growing understanding of the system. By digging into the tool's material properties, sorters began to develop new kinds of expertise based on understanding digital representations and algorithmic patterns. This allowed them to provide more useful explanations while reinforcing the value of their embodied expertise.

This new learning process quickly became entangled with day-to-day problem-solving. Supply chain managers started occasionally and periodically approaching the seed sorters when batches had ambiguous quality scores and they were unsure where to route the seed. In such situations, seed sorters tended to pull up the image records and walk through the visible patterns with managers, explaining how patterns such as uneven chlorophyll levels can signal variable germination. Managers asked follow-up questions about possible market destinations, and sorters related pattern interpretations to climatic conditions and logistics in response. These impromptu interactions became formalized over time, and seed sorters and supply chain managers began holding regular weekly meetings. During the meetings, sorters presented seed images, explained sorting patterns, and related them to specific climatic and logistical concerns. They also asked questions in return: 'Where is this batch going?' (Seed sorter) and 'What matters to this client?' (Seed sorter). These insights were important to seed sorters 'so [they] can know what to look for on images' (Seed sorter). As one seed sorter said, their new enhancing practice was fundamentally different from masking:

In the first stages of sorting advice, it was all human knowledge and we said this seed is bad, this seed is good. But now we say: yes, AI is taking over that. But still we need to say to [managers]: you need to trust the advice because it's also our advice. We know our machines, we know our analysis of the machines, we know how to use the information from AI to make advice for you. (Seed sorter)

The benefits of this shift for seed sorters' ability to convince managers became visible in practice. Supply chain managers began to rely on seed sorters not just for sorting results, but also for justification and evidence they could use in buyer negotiations. As one manager noted:

It helps a lot that now they can show us the images. Before, we just had the numbers. Now if I can see, for example, that most of the seed looks clean and uniform, it's easier to explain it to the [seed buyer]. (Supply chain manager)

This enhancing practice also became valuable to managers when responding to complaints. As one manager explained, in such situations managers began frequently calling seed sorters for more information:

Yeah, if we get a complaint, I just call [a seed sorter] and ask what happened. He explains to me what he sees and that the seed is good. He basically gives me proof to go back to the buyer and say the seed is good. (Supply chain manager)

In some cases, the value of seed sorters' explanations was in managers being able to dismiss buyer complaints entirely, by using the explanations to argue that 'it was a mistake of the buyer who didn't take good care of the seed, so we don't refund them' (Supply chain manager).

Encouraged by this new relevance, seed sorters worked on strengthening their authority by developing novel expertise relevant to enhancing explanations. For example, they initiated small research projects to analyse how image-based features correlated with field outcomes. By engaging in these activities, seed sorters were able to show more clearly the value of enhancing practice to their clients. For example, seed sorters organized internal training sessions where they explained to supply chain staff how they interpret image patterns and how they distinguish between model confidence and actual seed defects. One seed sorter explained how this strengthened their authority:

If there are complaints, we are responsible for validating the models. More importantly, we must teach others, like supply chain managers, how to read these images. They need to trust our decisions, so we teach them how we do it and why we are right. (Seed sorter)

This newly developed expertise was highly appreciated by managers. As sorters became more fluent in interpreting model outputs and connecting them to production realities, managers increasingly sought their input when facing uncertainty. One manager explained, 'If I have doubts, I go to them. They know how the model sees the seed and what that means for us in the field' (supply chain manager). By complementing the model with domain-specific insight, seed sorters ensured that their authority remained not only intact but also even expanded.

Recruitment Case

Traditional situation: Experts struggle to enact authority over clients by relying on their expertise. In the recruitment case, we observed how recruiters initially struggled to enact authority, in which senior managers – their key clients – often questioned selection decisions grounded in their expertise.

Traditionally, recruiters' role involved making selection decisions on job candidates: they were responsible for evaluating and screening candidates for job interviews with senior managers based on resumes, standardized test scores, and motivation letters. In doing so, they applied a set of predefined criteria based on a recruitment methodology and combined them with their judgment of relevant contextual factors, such as differences in the popularity of positions and countries candidates applied for. During a 10-minute briefing before every job interview, recruiters presented the selected candidates to senior managers using a slide deck. They explained why these candidates were suitable by summarizing their key impressions of their resumes and other documents, and advised senior managers on questions to ask during the job interviews.

Unlike in banking and biotechnology, clients often questioned recruiters' advice and selections of candidates despite their formal role in overseeing recruitment. For example, when facing job candidates who, in their eyes, performed poorly in interviews, managers complained to recruiters that they were 'wasting their time' (senior manager). One senior manager commented after a job interview: 'She was terrible, she just cannot handle the job! Why did she even pass the first selection round?' (senior manager). Senior managers' approval of job candidates was critical for recruiters, as they were evaluated on 'high interview pass rates' in their performance appraisals. The fact that senior managers questioned and dismissed their chosen candidates and advice, therefore, presented an important challenge to them. A recruiter shared:

Senior managers have very strong opinions after talking for just one hour to a candidate. It is tough to get through to them at this stage. (Recruiter)

Recruiters often struggled to convince senior managers of their chosen candidates, as their expertise was largely not recognized or appreciated as unique by senior managers. While they considered themselves 'recruitment experts' by bringing knowledge of recruitment methods, recruiters realized that senior managers benefited from business experience on what makes candidates succeed 'on the ground'. Senior managers believed that their hands-on experience gave them an edge over recruiters who lacked close involvement in their departments and teams, which would help them understand what qualities candidates need to possess. For example, a senior manager reflected on whether recruiters could reasonably judge if a candidate would succeed in the company:

Having a good conversation about candidates is important: Do I see this person working in my department or my team? It can be challenging to estimate through HR whether someone in sales, marketing, or supply will perform well. You need actual knowledge and experience from the business. (Senior manager)

As a result, recruiters struggled to establish authority over senior managers, as they lacked unique expertise in the domain that could give them an advantage over clients, as we observed in the banking and biotechnology cases.

Initial explaining practice: Showcasing AI outputs as a result of interpreting AI as an opportunity. The introduction of a new AI tool in candidate selection marked an important change in how experts sought to enact authority over clients. Although recruiters lost decision-making power to the technology, they viewed it as an opportunity to strengthen their position by relying on 'data-driven facts' rather than recruitment expertise, which their clients failed to appreciate. They explained chosen candidates to senior managers by openly referring to AI outputs, an explaining practice we refer to as *showcasing*.

Specifically, in mid-2018, HR management introduced an AI tool to evaluate and screen candidates for job interviews. The technology performed these evaluations without direct interference from recruiters by relying on machine learning algorithms that predict candidates' suitability based on inferred statistical patterns between personality

traits and employee performance. Recruiters received the algorithmic predictions in dashboards and simply needed to click a ‘filter out’ button to reject candidates based on a predefined threshold. They no longer had access to traditional selection instruments, such as resumes and motivation letters.

At the same time, recruiters remained in charge of presenting and explaining chosen candidates to senior managers, given their role in overseeing the selection process and quality of chosen candidates. To aid recruiters in explaining candidate selections, the AI tool generated several outputs, including personality scores and graphs with dominant personality traits underlying the prediction (see Appendix A). However, these outputs were experienced as largely opaque to recruiters, as they struggled to understand how the visual graphs and scores related to selection decisions:

I have no clue why these traits were chosen. Are they the strongest ones, the weakest ones? I really don't know how these words [in the word cloud] are put together. I understand the basics of how we work with AI, but it still seems abstract to me. (Recruiter)

Despite the disruption the technology presented to their responsibilities in candidate selection, recruiters viewed it as an opportunity to strengthen their credibility and demonstrate greater competence to managers. One recruiter reflected on how the AI outputs added weight to candidate selections, making it harder for managers to dismiss chosen candidates:

The reason we wanted to do that is really to try and make the recruitment process as fair and objective as possible. And to empower HR to come up with data points because, in the previous recruitment cycle, I saw many senior managers challenging presented candidates with vague arguments. (Recruiter)

Hence, in explaining selection decisions to managers, recruiters began actively *showcasing* AI outputs. Contrary to banking and biotechnology, where experts masked the role of technology to emphasize their own expertise, the practice of showcasing involved experts visibly and actively demonstrating their reliance on technology-generated results rather than their own expertise. Specifically, recruiters included the personality scores and graphs in their presentations to managers to illustrate how AI had been leading in candidate selection. For example, during a job interview, a recruiter would present a chosen candidate by citing and pointing out her AI outputs: ‘She has a super good AI score. And on the word cloud, you can see that she is assertive but modest as well’ (recruiter). By providing explanations that appeared ‘data-driven’ and free from subjective interpretation, recruiters aimed to convince managers of their decisions, knowing that managers often did not value their expert opinions. As shared by a recruiter:

I would say, like, that's interesting, but keep it as an observation and don't add my judgment to it. And then show the AI outputs as objectively as possible to senior managers. (Recruiter)

Senior managers were initially impressed by the data-driven explanations provided by recruiters regarding the chosen candidates. The emphasis on AI outputs stood in stark contrast to recruitment methodologies and criteria, which senior managers perceived as having a limited value in relation to their business experience. One senior manager expressed that the AI outputs made him feel that decisions about candidates were more legitimate and reliable:

With all of the data we now have in advance and the fact that you know someone has gone through previous robust rounds, you think: these people must have something. They must have strong qualities. (Senior manager)

Shifting explaining practice: Calibrating AI outputs by learning from clients. Although clients perceived the data-driven explanations as a subtle improvement, recruiters were confronted with a new challenge: the explanations were often seen as overly generic and lacking sufficient contextualization to be meaningful for senior managers. As managers increasingly voiced their concerns, recruiters gradually shifted their explaining practice – from simply showcasing the AI outputs to actively calibrating them in ways that aligned with managerial interests.

Starting in 2019, recruiters faced a growing number of questions and criticism during job interviews regarding their explanations of chosen candidates. Senior managers, despite being initially impressed by the AI outputs, became more critical as they engaged more deeply with the data and considered how to act upon it. They often interacted with the same team of recruiters during job interviews and began voicing their concerns about how the explanations lacked the clarity and contextual nuances needed to inform their interpretations and actions. For instance, they argued that the personality scores and visual graphs were not self-explanatory and required clear benchmarks and standards to be meaningful. As illustrated during a discussion between a manager and a recruiter:

Senior manager: So, a personality score of 59, is this good or bad?

Recruiter: A personality score of 59 is low.

Senior manager: But why is it low? When I see these things [personality scores], it is hard for me to understand without a reference point.

Given the increased emphasis on the AI outputs during interview briefings, senior managers began to challenge recruiters that the data presented had to be relevant, flagging this as an issue that required 'improvement':

People are very much looking at numbers. We are told that [a personality score] above 1000 is good, and then somewhere 1192 is, okay, that's a good candidate, 984, okay, that's not so good. Without any context. So that's something we need to improve on. (Senior manager)

Recruiters were deeply concerned by the criticism from senior managers, worrying that without clear and meaningful explanations, they would not receive the approval of

their clients needed to strengthen their authority. These managers were not only critical stakeholders but also individuals with whom recruiters regularly interacted during job interviews, intensifying pressures for recruiters to meet their expectations. As a result, they concluded that improving the AI outputs shown to managers to explain chosen candidates was a top priority for them: ‘The burning point for us is the data piece. What to show and what to explain to managers?’ (Recruiter).

Like biotechnology, recruiters were encouraged to revise their explaining practice as they learned about clients’ concerns and demands through repeated interactions. However, rather than investing in enhancing the AI outputs with their expertise, recruiters engaged in what we call *calibrating*: subtly adjusting, tweaking, and customizing the AI outputs to make decisions understandable and convincing for their clients. This practice focused on making explanations more contextualized and actionable for managers while maintaining the impression that these explanations remained neutral and free of expert judgment.

Starting at the end of 2019, recruiters engaged with AI developers in a 4-month exercise calibrating the AI outputs to make them more meaningful and self-explanatory for senior managers. This involved designing a new, automatically generated report for each candidate that represented their personality scores and graphs in a more contextual and actionable way (see Appendix A). Recruiters took an active role in shaping the report. For instance, they formulated explicit descriptions to accompany each personality score, thereby reducing the need for personal explanations. They also decided that candidates should be compared with internal benchmarks by visually plotting their traits against those of both high- and low-performing employees – groups that recruiters themselves deemed appropriate reference points. Additionally, they requested that the report automatically highlight those traits where candidates over- or underperformed and pull corresponding interview questions from a large internal database tied to those traits. Reflecting on these efforts, a recruiter explained that their aim was to make it easier for senior managers to understand both why a candidate was selected and which actions to take:

From senior managers, the request was to see candidates’ scores in relation to benchmarks for specific programs and our internal population. This is something we have built into the AI outputs. For example, we can now say: ‘AI shows that this candidate is underperforming on this specific trait. We recommend you ask this set of questions based on that’. (Recruiter)

Senior managers appreciated the calibrated AI outputs that are now featured in HR’s explanations of chosen candidates, noting that they provide a ‘clear reference point’ and ‘a good indication of where you should dig deeper’. They began incorporating them more actively in their questions and discussions on candidates. For example, a recruiter observed that managers started to formulate interview questions based on the AI outputs, interpreting this as a way that ‘their advice’ was taken up more strongly:

We really see that senior managers are using the [AI outputs on] lower traits to formulate interview questions. They refer to this a lot in the job interview, for example,

by asking: 'Ok, [candidate name], AI shows us you score low on "openness to experience". Can you tell us about the last time you put yourself out of your comfort zone?' (Recruiter)

Senior managers also began to refer to the AI outputs to resolve dilemmas regarding candidates and support their decisions on job candidates. As illustrated during a job interview:

She was the love or hate of the job interview. When she spoke, she demonstrated remarkable insight and intelligence. At the same time, I doubt her motivation. But man, I understand it at her age. That was my thinking, and that's why I would hire her. And, importantly, her AI score was the highest of all candidates. She scored more than a 1000! And she had a wide range of skills and traits. That's why I think it's worth giving her a shot. (Senior manager)

Recruiters were thrilled with these changes in how managers were using the AI outputs, viewing them as evidence that they were better able to convince managers of chosen candidates and steer their actions in stronger ways than before. A recruiter shared how he felt that the AI outputs 'empowered' them through the promise of objectivity they provided, in which he and his colleagues were inclined to place even more emphasis on the AI outputs when briefing managers on candidates:

Some of our leaders are difficult to convince or change their minds. The AI outputs finally empower the HR team by having an objective data point to refer to [in discussions with managers]. We now see recruiters referring back to the scores and graphs even more in the interview briefings. (Recruiter)

In the short run, the practice of calibrating, therefore, enabled recruiters to lure managers into accepting and acting upon AI decisions, thereby strengthening their authority. Yet, by presenting their involvement as minimal and the explanations as products of the technology alone, they missed out on opportunities to build recognition for their expertise in the long run. Although our fieldwork ended before such longer term consequences could unfold, our comparative analysis suggests that such practices may ultimately erode the sustainability of recruiters' authority, as their expert contributions were being absorbed into outputs that appear neutral and self-sufficient. Enchanted by the promise of objectivity, recruiters themselves appeared largely unaware of these pitfalls, focusing instead on how 'visibility is automatically provided' (recruiter). As another recruiter summarized: 'Our dream is to plug in that extra visibility we created and let the tool speak for itself'.

CROSS-CASE ANALYSIS

Our cross-case analysis reveals commonalities and differences among the cases we studied. In all cases, the use of predictive AI tools for decision-making disrupted how experts traditionally enacted authority over clients. Instead of relying on their expertise to gain

clients' approval, experts found themselves confronted with a new technology that generated decisions and presented outputs that were at odds with experts' judgment. Still, experts had to explain these AI decisions to their clients. To reconstruct their authority over clients, experts engaged in what we call *explaining practices*; activities aimed at making clients understand and accept AI decisions. These practices were shaped by two relational conditions – whether clients recognize the expert's knowledge as unique, and whether interactions with clients provide rich opportunities for experts to learn about clients' evolving needs – with important consequences for whether expert authority was ultimately strengthened. Figure 1 presents our model of how expert authority is enacted in light of an AI disruption.

Our model starts by showing how experts enact different explaining practices depending on whether they benefit from expertise that is recognized as *unique* by their clients. Our analysis suggests that expert groups who traditionally earn clients' approval through clients recognizing their unique expertise interpret AI as threatening their authority by undermining their judgment. These groups strive to defend their expertise as a source of authority by masking AI outputs and relying on their prior experience and judgment when explaining AI decisions. In contrast, expert groups who struggle to gain clients' approval through their expertise may interpret the technology as an opportunity by offering a 'data-driven' alternative to their judgment, which is typically questioned. We found that these experts explain AI decisions by actively showcasing AI outputs, rather than emphasizing their own expertise.

Our analysis highlights that while experts carefully consider how clients recognize their expertise, their explaining practices often fall short of meeting clients' evolving needs, as they tend to remain generic and decontextualized. This issue of decontextualization occurs as, on the one hand, experts continue to rely on traditional decision rules that have lost relevance with AI decision-making procedures. For instance, loan consultants explained loan rejections based on familiar expert criteria; yet, these criteria no longer reflected how individual decisions were actually made by the AI system. On the other hand, experts may overly rely on AI outputs that lack contextual




Figure 1. A model of how experts aim to enact authority over clients in light of AI disruption

grounding. For instance, recruiters presented the AI scores and graphs with minimal interpretation and translation in an attempt to appear ‘objective’, but offered little guidance to help managers understand how these numbers related to standards for roles relevant in concrete contexts.

Expert groups differ in their opportunities for learning about and addressing the issue of decontextualization, resulting in some groups shifting their explaining practices while others stagnate. Specifically, experts who continuously and closely engaged with clients were directly confronted with the shortcomings of generic, decontextualized explanations and felt pressure to revise their practices. In contrast, experts who only sparsely interacted with clients, as in the case of one-off exchanges in banking, had fewer opportunities to learn about client needs and thus continued to rely on practices that failed to provide detailed and contextualized insights. Importantly, among those who revised their practices, clients’ recognition of expertise still shaped how adaptations were made. When clients recognized the value of domain-specific knowledge, experts responded by enhancing AI outputs with newly developed expertise. Conversely, when clients disregarded their expertise, experts withdrew further from the explanatory process, opting instead to calibrate AI outputs so that explanations were ‘automatically provided’.

The ability of experts to learn about and act upon their decontextualized explanations had important implications for their authority over clients. We find that expert groups who managed to better tailor their explanations to clients’ needs lured clients into following issued decisions and advice, thereby strengthening their authority. For example, as clients sought detailed, context-specific explanations of AI decisions, seed sorters investigated AI’s material properties, and this ongoing exchange led to the development of new expertise. As a result, supply chain managers increasingly relied on the sorters’ explanations when acting on AI-generated decisions. In contrast, expert groups who lacked opportunities for learning about the need for contextualized explanations failed to meet clients’ interests. As a result, clients remained unconvinced by the explanations provided and were more likely to seek alternatives, as observed in the banking case. Overall, our cross-case analysis thus highlights that when experts lose decision-making control to AI, their authority is not simply diminished but reconfigured through the ways they craft explanations for their clients – shaped by relational conditions of clients’ recognition of expertise and learning.

THEORETICAL IMPLICATIONS

The relational perspective on how expert authority is enacted in light of AI decisions put forward in this study has important implications for the literature on professions and expert occupations, as well as a growing body of research on AI in management.

First, our study contributes to the literature on professions and expert occupations by revealing alternative pathways through which experts enact authority when facing technological disruptions. Moving beyond studies that locate authority primarily in credentials (e.g., Greenwood et al., 2002; Scott, 2008) or abstract knowledge (e.g., Abbott, 1988; Freidson, 1985; Hughes, 1958), we offer a relational perspective that

foregrounds experts' daily interactions with immediate clients as central in understanding how experts attempt to restore their authority in the face of AI. Our findings show that formal roles and abstract knowledge do not shield experts from the challenges posed by AI; instead, they had to resort to alternative means to strengthen their authority. Our model identifies a set of explaining practices through which experts sought to maintain their relevance, with the effectiveness of these practices hinging on experts' ability to relate and respond to the needs of their clients. This study thus emphasizes that authority in the age of AI cannot be evaluated solely through formal roles or expertise, but must also be understood in relation to the explanations experts craft for their audiences.

Second, we contribute to workplace studies on how experts navigate AI disruptions by shifting the focus from preserving expertise to reconfiguring it in response to client needs. Prior studies have largely highlighted that experts try to hold on to their expertise when AI tools enter their work to emphasize the continued relevance of their position. For example, experts may protect their positions by substituting AI predictions with their own expertise (e.g., Anthony, 2021; Lebovitz et al., 2022; Waardenburg et al., 2022) or by symbolically complying with AI outputs to maintain their traditional expert practices (e.g., Pachidi et al., 2021). Our comparative field study nuances this view by showing that experts leverage their expertise differently depending on whether clients recognize it as unique and valuable: while groups whose clients traditionally value their expertise mask or enhance AI outputs to signal the continued relevance of their expert position, others move away from expert rules in favour of seemingly 'objective' data patterns to strengthen their claims. These practices carry real consequences for how expertise is demonstrated, performed, and developed amid technological change. This study thereby challenges the prevailing view that expertise is merely preserved in the face of AI and instead highlights how it is mobilized differently through client-expert dynamics.

Third, our study also contributes to organizational and management research on how AI shapes the development of new expertise. Prior studies have shown that domain experts can cultivate new forms of expertise through close engagement with AI systems and their designers, including practices of interrogating AI outputs in relation to their own knowledge claims (Lebovitz et al., 2022), engaging in mutual learning with AI developers (Van den Broek et al., 2021), and recreating boundaries to integrate different expertise from developers (Faulconbridge et al., 2023). We extend this literature by highlighting an alternative pathway for expertise development that centres on making AI outputs meaningful through sustained interactions with outside audiences. In our biotechnology case, experts learned through repeated feedback from clients who questioned their generic explanations. This prompted experts to explore new ways of engaging with the system's features and outputs to refine their interpretations and make outputs meaningful in client-specific contexts. We thus demonstrate that opportunities for developing new expertise arise not only through direct engagement with AI systems and their creators but also through the demands of those who face and evaluate AI-generated decisions.

Finally, this study advances organizational studies on AI explainability by deepening our understanding of how experts approach the issue of decontextualization

surrounding AI explanations. Nascent research has recognized that technical approaches – such as feature maps or simplified decision rules – often produce explanations detached from the situated realities of their audiences, highlighting the need to tailor explanations to particular groups (Hafermalz and Huysman, 2022; Pakarinen and Huisman, 2023; Waardenburg et al., 2022). Our study extends this work by demonstrating that decontextualization arises not only when experts lean too heavily on technical outputs (as in recruitment) but also when they continue to communicate traditional decision rules that have lost relevance under AI decision-making procedures (as in banking and biotechnology). Critically, our comparative analysis highlights that experts do not respond to this issue uniformly but rather diverge based on the kinds of relationships they hold with their clients. Experts embedded in ongoing exchanges are repeatedly confronted with the shortcomings of generic explanations and pressured to contextualize AI outputs through practices of enhancing or calibrating. By contrast, those with sparse interactions encounter fewer opportunities to learn about these shortcomings and remain locked into decontextualized practices. This study thereby demonstrates that different forms of expert-client relationships shape not only the need for but also the ability to contextualize AI explanations in practice.

PRACTICAL IMPLICATIONS

Our findings yield practical insights for managers and practitioners who implement and use AI in decision-making. First, we emphasize that the pursuit of ‘explainable AI’ (XAI) by developers and regulators (e.g., European Union, 2024; Gunning et al., 2019) risks reducing explanations to merely a technical property of AI tools. This approach typically centres on algorithmic transparency and post-hoc interpretability tools, which are assumed to make AI decisions understandable and acceptable to end users. However, as these technical explanations often fail to resonate with client needs unless embedded in socially meaningful interactions, our findings emphasize the importance for managers, experts, and AI developers to refrain from viewing explainability as a static feature and instead treat it as a dynamic, relational process co-constructed with clients in practice. For instance, this can be achieved by encouraging open dialogue during client meetings, where explanations are continuously adapted based on clients’ questions, prior knowledge, and concerns.

Second, for experts tasked with presenting AI decisions, it is now more important than ever to invest efforts in understanding clients’ needs and actively tailoring explanations. This includes learning how clients use the information, asking what explanations they find most useful, and adjusting presentation styles accordingly by, for instance, visual checks to highlight relevant details, preparing simple summaries with familiar markers of quality, or using terms clients already know. Such efforts help integrate AI decisions into existing workflows and maintain clients’ trust. Managers may play a crucial role in facilitating opportunities for rich interaction and feedback between experts and clients. For instance, managers may organize learning sessions where clients and experts can openly share challenges related to the provided explanations or integrate designated feedback moments during expert-client encounters. These interactions are crucial to ensure that experts derive compelling and

informative explanations that connect to clients' needs, while also articulating the potential and pitfalls of AI technologies. Recent examples, such as the Dutch tax scandal (Toh, 2020), serve as a cautionary tale regarding the ethical risks that arise when experts merely 'hide' behind the AI tool without taking responsibility. Therefore, managers and experts responsible for implementing and using these technologies need to acknowledge the importance of involving clients in learning about the potential risks of AI tools as early as possible.

BOUNDARY CONDITIONS AND AVENUES FOR FUTURE RESEARCH

Our findings need to be considered in light of certain boundary conditions that offer new directions for future research. First, our study particularly focuses on how organizations use AI to delegate high-stakes tasks. These tasks are central to expert work and often particularly sensitive to technological disruption as they pose a significant vulnerability for experts (e.g., Faulconbridge et al., 2023; Selenko et al., 2022). While this focus allows us to highlight the important role of clients in shaping experts' responses to AI tools, prior research leads us to expect that experts may respond differently depending on the task at stake or the expert's role identity (e.g., Faulconbridge et al., 2023; Selenko et al., 2022). Thus, we encourage future research to explore how expert authority over clients is shaped in various contexts, for diverse tasks, and across different expert groups.

Second, our case focuses on a specific class of AI: predictive AI based on supervised machine learning. As new forms of AI, such as generative AI, enter organizations, new dynamics may emerge in the relationships we observed. Generative AI technologies may enable experts to tailor explanations more carefully to client needs and contexts through techniques like prompting (Mayer et al., 2025a, 2025b; Retkowsky et al., 2024). At the same time, these technologies may reduce opportunities for experts to play a role in this process by promising to 'explain' AI outputs directly to clients. For instance, advanced Large Language Models (LLMs) offer new means for clients as laypeople to interpret AI outputs themselves. Thus, an opportunity lies ahead for future research to explore how explaining practices evolve as both experts and clients interact with generative AI tools and with what consequences for expert authority.

Finally, while our study builds on longitudinal data, future research could further examine how explaining practices unfold over time and across different audiences. Specifically, our findings suggest that practices that strengthen expert authority in the short term may be less sustainable in the long run. In the recruitment case, for example, recruiters increasingly led senior managers into accepting and acting upon AI decisions, yet largely concealed their own role in the process. Unlike seed sorters, who invested in developing and demonstrating new, complementary expertise, recruiters' emphasis on 'impartial' explanations obscured their involvement. These experts may face a dilemma in the long run that warrants closer attention: on the one hand, experts may persuade clients to accept AI decisions, but at the same time, they lose the ability to signal and justify their involvement, potentially undermining their long-term authority. Moreover, our banking case suggests that experts may face different audiences

over time that shape these dynamics. Groups like higher management, technical staff, or external consultants can also trigger experts to revisit their explaining practices, raising new pressures on how authority is negotiated. This calls for more longitudinal research on the durability of different explaining practices and how varied audiences shape the evolution of these practices.

CONCLUSION

In a world where AI tools are increasingly entering the workplace of experts, our study uncovers the relational dynamics that surround expert authority over clients when they are confronted with AI decisions. Drawing on a comparative field study, we show that when experts lose decision-making control to AI, their authority is not simply diminished but reconfigured through the ways they craft explanations for their clients – shaped by the relational conditions of clients' recognition of expertise and learning. This study thereby offers new insights into the complex interplay between expert-client relationships, expert authority, and explaining practices.

ACKNOWLEDGMENTS

We are grateful to our Special Issue editor Stella Pachidi, editor Kristina Potočnik, and the three anonymous reviewers for their very constructive and valuable guidance, which greatly helped us in improving our work. We would also like to thank Daisy Chung, Anastasia Sergeeva, Reza Baygi, and Marleen Huysman for their helpful feedback on earlier drafts of this paper. Further, this research has benefited greatly from discussions with participants at the International Conference on Information Systems, the European Group for Organizational Studies, and the Annual Meeting of the Academy of Management. This work was partly funded by the Marianne and Marcus Wallenberg Foundation (grant 2021.0133). Open Access funding enabled and organized by Projekt DEAL.

REFERENCES

Abbott, A. (1988). *The System of Professions: An Essay on the Expert Division of Labor*. Chicago, IL: University of Chicago Press.

Anteby, M., Chan, C. K. and DiBenigno, J. (2016). 'Three lenses on occupations and professions in organizations: Becoming, doing, and relating'. *Academy of Management Annals*, **10**, 183–244.

Anthony, C. (2018). 'To question or accept? How status differences influence responses to new epistemic technologies in knowledge work'. *Academy of Management Review*, **43**, 661–79.

Anthony, C. (2021). 'When knowledge work and analytical technologies collide: The practices and consequences of black boxing algorithmic technologies'. *Administrative Science Quarterly*, **66**, 1173–212.

Anthony, C., Bechky, B. A. and Fayard, A.-L. (2023). '"Collaborating" with AI: Taking a system view to explore the future of work'. *Organization Science*, **34**, 1672–94.

Bechky, B. A. and Okhuysen, G. A. (2011). 'Expecting the unexpected? How SWAT officers and film crews handle surprises'. *Academy of Management Journal*, **54**, 239–61.

Bechky, B. A. and O'Mahony, S. (2015). 'Leveraging comparative field data for theory generation'. In Elsbach K. D. and Kramer R. (Eds), *Handbook of Qualitative Organizational Research*. New York: Routledge, 200–8.

Berente, N., Gu, B., Recker, J. and Santhanam, R. (2021). 'Managing artificial intelligence'. *MIS Quarterly*, **45**, 1433–50.

Burrell, J. (2016). 'How the machine 'thinks': Understanding opacity in machine learning algorithms'. *Big Data & Society*, **3**(1), 1–12.

Carr-Saunders, A. M. and Wilson, P. A. (1933). *The Professions*. Oxford: Clarendon Press.

Cecez-Kecmanovic, D. (2025). 'Ethics in the world of automated algorithmic decision-making—a posthumanist perspective'. *Information and Organization*, **35**, 100587.

Charmaz, K. (2006). *Constructing Grounded Theory: A Practical Guide Through Qualitative Analysis*. London: Sage.

Christin, A. (2020). 'The ethnographer and the algorithm: Beyond the black box'. *Theory and Society*, **49**, 897–918.

DiBenigno, J. (2018). 'Anchored personalization in managing goal conflict between professional groups: The case of US Army mental health care'. *Administrative Science Quarterly*, **63**, 526–69.

DiBenigno, J. (2020). 'Rapid rationality: How peripheral experts build a foundation for influence with line managers'. *Administrative Science Quarterly*, **65**, 20–60.

Eisenhardt, K. M. (1989). 'Building theories from case study research'. *Academy of Management Review*, **14**(4), 532–50.

Emirbayer, M. (1997). 'Manifesto for a relational sociology'. *American Journal of Sociology*, **103**, 281–317.

European Union (2024). 'Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 March 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts'. *Official Journal of the European Union*, **259**, 1–125. Available at: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj> (accessed 29 February 2025).

Eyal, G. (2019). *The Crisis of Expertise*. New Jersey: John Wiley & Sons.

Faraj, S., Pachidi, S. and Sayegh, K. (2018). 'Working and organizing in the age of the learning algorithm'. *Information and Organization*, **28**, 62–70.

Faulconbridge, J., Sarwar, A. and Spring, M. (2023). 'How professionals adapt to artificial intelligence: The role of intertwined boundary work'. *Journal of Management Studies*, **62**, 1991–2024.

Freidson, E. (1985). 'The reorganization of the medical profession'. *Medical Care Review*, **42**, 11–35.

Greenwood, R., Suddaby, R. and Hinings, C. R. (2002). 'Theorizing change: The role of professional associations in the transformation of institutionalized fields'. *Academy of Management Journal*, **45**, 58–80.

Gunning, D., Stefk, M., Choi, J., Miller, T., Stumpf, S. and Yang, G.-Z. (2019). 'XAI – explainable artificial intelligence'. *Science Robotics*, **4**, eaay7120.

Hafermalz, E. and Huysman, M. (2022). 'Please explain: Key questions for explainable AI research from an organizational perspective'. *Morals and Machines*, **1**, 10–23.

Huang, M.-H., Rust, R. and Maksimovic, V. (2019). 'The feeling economy: Managing in the next generation of artificial intelligence (AI)'. *California Management Review*, **61**, 43–65.

Hughes, E. C. (1958). *Men and Their Work*. London: Collier-Macmillan Limited.

Huising, R. (2015). 'To hive or to hold? Producing professional authority through scut work'. *Administrative Science Quarterly*, **60**, 263–99.

Huising, R. (2023). 'Professional authority'. In Dans, Eyal, G. and Medvetz, T. (Eds), *The Oxford Handbook of Expertise and Democratic Politics*. Cambridge: Oxford University Press, 453–69.

Kaynak, E. and Barley, S. R. (2019). 'Shaping the political environment: An ethnography of public affairs professionals at work'. *Work and Occupations*, **46**, 265–306.

Klutz, D. N. and Mulligan, D. K. (2019). 'Automated decision support technologies and the legal profession'. *Berkeley Technology Law Journal*, **34**, 853–90.

Langley, A. (1999). 'Strategies for theorizing from process data'. *Academy of Management Review*, **24**(4), 691–710.

Larson, M. (1979). *The Rise of Professionalism: A Sociological Analysis*. Berkeley, CA: University of California Press.

Lebovitz, S., Lifshitz-Assaf, H. and 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**, 126–48.

Levina, N. and Vaast, E. (2005). 'The emergence of boundary spanning competence in practice—Implications for implementation and use of information systems'. *MIS Quarterly*, **29**(2), 335–63.

Li, C., Wang, H., Jiang, S. and Gu, B. (2024). 'The effect of AI-enabled credit scoring on financial inclusion: Evidence from one million underserved population'. *MIS Quarterly*, **48**, 1803–34.

Lund, S., Manyika, J., Segel, L., Dua, A., Hancock, B., Rutherford, S. and Macon, B. (2019). *The Future of Work in America: People and Places, Today and Tomorrow*. New York: McKinsey Global Institute.

Mayer, A.-S., Kostis, A., Strich, F. and Holmström, J. (2025a). 'Shifting dynamics: How generative AI as a boundary resource reshapes digital platform governance'. *Journal of Management Information Systems*, **42**, 400–30.

Mayer, A.-S., Baygi, R. M. and Buwalda, R. (2025b). *Generation AI: Job Crafting by Entry-Level Professionals in the Age of Generative AI*. Business & Information Systems Engineering.

Mayer, A.-S., Strich, F. and Fiedler, M. (2020). 'Unintended consequences of introducing AI systems for decision making'. *MIS Quarterly Executive*, **19**, 239–57.

Mukherjee, S. and Thomas, C. D. (2023). 'Feeling the rule management and relational authority: Fostering patient compliance in palliative care consultations'. *Organization Studies*, **44**, 451–72.

O'Mahony, S. and Bechky, B. A. (2008). 'Boundary organizations: Enabling collaboration among unexpected allies'. *Administrative Science Quarterly*, **53**, 422–59.

Pachidi, S., Berends, H., Faraj, S. and Huysman, M. (2021). 'Make way for the algorithms: Symbolic actions and change in a regime of knowing'. *Organization Science*, **32**, 18–41.

Pakarinen, P. and Huisng, R. (2025). 'Relational expertise: What machines can't know'. *Journal of Management Studies*, **62**(5), 2053–82.

Prasad, P. (1993). 'Symbolic processes in the implementation of technological change: A symbolic interactionist study of work computerization'. *Academy of Management Journal*, **36**, 1400–29.

Retkowsky, J., Hafermalz, E. and Huysman, M. (2024). 'Managing a ChatGPT-empowered workforce: Understanding its affordances and side effects'. *Business Horizons*, **67**(5), 511–23.

Scott, W. R. (2008). 'Lords of the dance: Professionals as institutional agents'. *Organization Studies*, **29**, 219–38.

Selenko, E., Bankins, S., Shoss, M., Warburton, J. and Restubog, S. L. D. (2022). 'Artificial intelligence and the future of work: A functional-identity perspective'. *Current Directions in Psychological Science*, **31**, 272–9.

Staudenmayer, N., Tyre, M. and Perlow, L. (2002). 'Time to change: Temporal shifts as enablers of organizational change'. *Organization Science*, **13**, 583–97.

Strich, F., Mayer, A.-S. and Fiedler, M. (2021). 'What do I do in a world of artificial intelligence? Investigating the impact of substitutive decision-making AI systems on employees' professional role identity'. *Journal of the Association for Information Systems*, **22**, 9.

Toh, A. (2020). Dutch Ruling a Victory for Rights of the Poor. Human Rights Watch. Available at <https://www.hrw.org/news/2020/02/06/dutch-ruling-victory-rights-poor> (accessed 15 October 2024).

Van den Brock, E., Sergeeva, A. and Huysman, M. (2021). 'When the machine meets the expert: an ethnography of developing AI for hiring'. *MIS Quarterly*, **45**, 1557–80.

Van den Brock, E., Sergeeva, A. and Huysman, M. (2025). 'Is there fairness in AI?'. *Journal of Management Studies*. <https://onlinelibrary.wiley.com/doi/full/10.1111/joms.13276>

Waardenburg, L., Huysman, M. and Sergeeva, A. V. (2022). 'In the land of the blind, the one-eyed man is king: Knowledge brokerage in the age of learning algorithms'. *Organization Science*, **33**, 59–82.

Yin, R. K. (1992). 'The case study method as a tool for doing evaluation'. *Current Sociology*, **40**(1), 121–37.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at the publisher's web site.