

Does AI Actually Help? How Employees Judge the Effectiveness of AI Tools in the Workplace

Xin Yue

Universidad Europea del Atlántico, Santander, Cantabria, Spain

ABSTRACT

Organisations around the world are spending heavily on artificial intelligence tools, trusting that those investments will make their employees faster, smarter, and more productive. But the people actually sitting at those desks tell a more complicated story. This paper critically examines how employees measure the effectiveness of AI tools they use in their day-to-day organisational roles, and why trust in those tools is far from automatic. Drawing on the technology acceptance literature, automation trust research, and documented workplace cases, the paper shows that employee judgments of AI effectiveness are shaped by perceived usefulness, transparency of outputs, past errors, and the degree to which a tool genuinely reduces cognitive load rather than adding to it. The paper also identifies the conditions under which AI tools erode rather than build employee trust, and it concludes with a set of practical advisories for workplace managers who want to close the gap between what AI promises and what it actually delivers.

Keywords: AI tools, employee perception, technology acceptance, algorithm aversion, workplace trust, AI effectiveness

INTRODUCTION

Few workplace conversations in recent years have generated as much excitement, as much anxiety, and as much confusion as the question of what artificial intelligence can actually do for employees. Organisations have adopted AI tools for tasks ranging from drafting emails and summarising meeting notes to screening job applications, predicting customer churn, and assisting doctors in reading medical images. The pitch, almost universally, is the same: AI will make your people more effective, your processes faster, and your decisions smarter.

The evidence, when examined closely, is considerably more uneven. Some employees report genuine relief from repetitive work. Others describe AI tools that produce outputs they cannot fully understand, cannot reliably verify, and do not fully trust. A growing number find that checking, correcting, or explaining AI-generated outputs has become a task in itself, adding steps to workflows that were supposed to be simplified. This is not a fringe experience. Surveys consistently show that large proportions of employees express scepticism about the AI tools they are required to use, even when those tools come with confident organisational endorsements.

What determines whether employees see an AI tool as genuinely helpful or as an additional complication? That question sits at the heart of this paper. Effectiveness, as this paper frames it, is not a property of the technology alone. It is a judgment that employees make, shaped by their direct experience with the tool, their understanding of how it works, their memory of when it has gone wrong, and their sense of whether the tool respects their professional knowledge or undermines it.

This paper draws on the established literature on technology acceptance and automation trust, examines documented cases from real workplaces where AI adoption has produced trust deficits rather than

confidence, and closes with concrete advisories for workplace managers who want to make AI tool adoption go better than it often does. Throughout, the focus stays on the employee's point of view, because that is the point of view that ultimately determines whether an AI investment succeeds or sits unused on a corporate server.

WHAT THE LITERATURE TELLS US ABOUT EMPLOYEES AND TECHNOLOGY

The question of how employees decide whether a new technology is worth using has been studied systematically since the 1980s. Davis (1989) introduced the Technology Acceptance Model (TAM), which identified two factors that consistently predict whether people adopt a new tool: perceived usefulness and perceived ease of use. Perceived usefulness refers to the degree to which a person believes using the system will improve their job performance. Perceived ease of use refers to how little mental effort the person expects to expend. These two dimensions have held up remarkably well across decades of technology adoption research, and they are directly applicable to how employees approach AI tools today.

Venkatesh et al. (2003) built on this foundation with the Unified Theory of Acceptance and Use of Technology (UTAUT), which added social influence and facilitating conditions to the picture. Social influence captures the fact that employees are more likely to adopt a tool if colleagues they respect are seen using it, and facilitating conditions address whether the organisation is providing the infrastructure and support that make adoption practical. Both factors matter when it comes to AI. An employee whose trusted senior colleague treats an AI tool with scepticism is unlikely to embrace it enthusiastically, regardless of what the vendor's brochure says.

The dimension of trust adds another layer. Lee and See (2004) offered a foundational account of how people form trust relationships with automated systems. They argued that trust in automation is not binary but calibrated: people adjust their reliance on automated systems based on their sense of the system's competence, the system's consistency over time, and the degree to which they can predict when the system will fail. This calibration process is healthy when it functions well. When employees over-trust a system, they stop scrutinising its outputs, which can allow errors to pass undetected. When they under-trust it, they stop using it at all, which wastes the investment and frustrates the people who mandated the adoption.

Mayer et al. (1995) provided a widely used framework for organisational trust that is equally applicable here. They defined trust as a willingness to be vulnerable to another party's actions, based on assessments of that party's ability, benevolence, and integrity. Translated to AI, employees assess whether a tool is capable of doing what it claims, whether it seems designed with their interests in mind, and whether it is honest about its limitations. An AI tool that confidently produces wrong answers fails on all three dimensions at once.

One of the most documented and counterintuitive phenomena in this space is algorithm aversion. Dietvorst et al. (2015) showed in a series of experiments that people become significantly less willing to rely on an algorithm after they have seen it make an error, even when the algorithm's overall accuracy is higher than human judgment. This is not what happens with human experts. When a human colleague makes a mistake, we tend to make allowances. When an algorithm makes a mistake, we question the entire system. The implications for workplace AI adoption are serious. A single visible, embarrassing error by an AI tool can undermine months of successful performance in the eyes of the employees who witnessed it.

Logg et al. (2019) found evidence of the opposite tendency in some populations, which they called algorithm appreciation: a preference for algorithmic over human judgment, particularly among people who have less domain expertise. This suggests that employee responses to AI tools are not uniform. Experienced professionals with deep domain knowledge are more likely to scrutinise and resist AI recommendations. Less experienced employees may accept AI outputs more readily, sometimes

uncritically. Both patterns carry risks, and neither represents the thoughtful, calibrated relationship with AI that actually produces good outcomes.

Parasuraman and Riley (1997) framed these risks in terms of misuse, disuse, and abuse of automation. Misuse occurs when employees rely on automated systems inappropriately, beyond the system's competence. Disuse occurs when employees ignore or reject systems that could genuinely help them. Abuse occurs when automation is used in ways it was not designed for. All three patterns are visible in contemporary AI adoption, and all three represent failures of the relationship between employee and tool, failures that better management could prevent.

HOW EMPLOYEES ACTUALLY MEASURE AI EFFECTIVENESS

When employees assess whether an AI tool is working for them, they tend not to conduct the kind of systematic evaluation that an IT department might. They apply a set of practical, experiential tests, often without being fully aware they are doing so.

The most immediate test is speed. Does this tool save me time? Employees who find that an AI tool takes longer to prompt, correct, and verify than it would have taken them to do the task themselves tend to abandon it quickly, regardless of what the theoretical time savings are supposed to be. Speed matters most for routine, repetitive tasks where the value proposition is most obvious. It matters less for complex professional judgments, where employees often prefer to think things through themselves.

The second test is accuracy. Does this tool get things right? This is where algorithm aversion becomes practically significant. Employees who catch an AI tool producing a confident but incorrect output do not forget it. They may continue using the tool, but they will verify its outputs more carefully, which adds work. If the errors continue, they will stop using the tool altogether. Davenport and Ronanki (2018) noted that AI projects in organisations frequently underperform because the organisations involved failed to anticipate how much error-checking would fall on the employees using the tools.

The third test is transparency. Do I understand why this tool is telling me what it is telling me? Employees who cannot see how an AI tool reached its recommendation face a specific difficulty: they cannot tell the difference between a case where the AI is right and they are wrong, and a case where the AI is wrong and their own instinct is right. Without transparency, the only rational strategy is either to trust the tool completely or to ignore it completely. Neither is what organisations want.

The fourth test is what might be called task fit. Is this tool actually designed for the kind of work I do? Employees are perceptive about the difference between a tool that was built for their specific context and one that is a general-purpose product being applied to a role it was not really designed for. Task fit problems are common in AI adoption because organisations often purchase tools based on category rather than specificity. A language model that works well for generating marketing copy may perform poorly when asked to help a compliance officer draft regulatory correspondence, and the compliance officer will notice immediately.

The fifth and perhaps most important test is cognitive load. Does this tool make my job feel lighter or heavier? Brynjolfsson and McAfee (2014) pointed out that new technologies do not automatically reduce cognitive demands; they often shift those demands to new locations. An AI tool that generates ten candidate responses to a customer email has technically done work, but if the employee now has to read and evaluate all ten before deciding which one is acceptable, the total cognitive work may be greater than it was before. Employees who experience this pattern reliably report dissatisfaction with the tool, even if they struggle to articulate exactly why.

TRUST DEFICITS: WHEN AI MAKES THINGS WORSE

Trust deficits in workplace AI adoption are not theoretical possibilities. They have been documented in enough organisational settings, across enough industries, to constitute a recognisable pattern. The following cases illustrate how that pattern plays out.

1. Amazon's Recruiting Algorithm

In 2018, Reuters reported that Amazon had developed and subsequently abandoned an AI recruiting tool that had been trained on a decade of historical hiring data. The system learned from the data that the company had historically hired more men than women in technical roles, and it began downgrading resumes that included signals of female identity, such as attendance at all-women colleges or membership in women's professional organisations. The engineers who discovered this tried to correct it, but the company ultimately concluded that the problem could not be reliably fixed and discontinued the project.

The trust deficit this created had several layers. The employees involved in the project lost confidence in the system's outputs. The candidates who might have been screened out had no knowledge of or recourse against the process. And the wider organisation was confronted with the uncomfortable realisation that a system designed to remove human bias had, in fact, encoded and amplified it. The lesson the literature would draw from this is precisely the one Mayer et al. (1995) anticipated: a system that fails on integrity, even if its capability is not in doubt, loses the trust of the people it is supposed to serve.

2. IBM Watson Health and Clinical Oncology

IBM's Watson for Oncology was rolled out at a number of major hospitals in the mid-2010s as a tool that could assist oncologists in selecting cancer treatments. The promise was compelling. Watson had been trained on an enormous body of medical literature and case data, and IBM marketed it as capable of identifying treatment options that human clinicians might miss.

The actual clinical experience was considerably less encouraging. Reporting by STAT News in 2017 and 2018 revealed that clinicians at institutions including Memorial Sloan Kettering and major hospitals in India and South Korea had found Watson making treatment recommendations that senior oncologists considered unsafe or medically unsound. Doctors reported that the system sometimes recommended treatments based on a small number of hypothetical cases rather than real patient data, and that it was difficult to understand how the recommendations were being generated. Clinicians who could not understand the reasoning behind a recommendation were understandably reluctant to act on it when the stakes were a patient's life.

This case illustrates the transparency problem at its most consequential. The algorithm aversion effect documented by Dietvorst et al. (2015) was visible in its clinical form: once oncologists had seen the system produce recommendations they believed were wrong, they did not just become cautious about those particular cases. They became broadly sceptical of the tool. Some stopped consulting it altogether. The tool was not worthless, but its adoption had been managed in a way that allowed its failures to be more visible than its successes, and the trust deficit that followed proved very difficult to reverse.

3. COMPAS and Criminal Justice

The COMPAS algorithm, used by courts in the United States to assess the likelihood that a defendant would reoffend, became the subject of a landmark investigation by ProPublica in 2016. ProPublica's analysis found that the algorithm was incorrectly flagging Black defendants as future criminals at roughly twice the rate it was doing so for white defendants, while simultaneously mislabelling white defendants who did go on to reoffend as lower risk.

The trust deficit this created extended in two directions. Criminal justice professionals who had been using the tool were confronted with evidence that the risk scores they had been incorporating into sentencing recommendations were racially biased. Some responded by distancing themselves from the tool; others continued using it while trying to mentally correct for the identified bias, which is itself a questionable practice. Defendants and their legal representatives, meanwhile, had been subject to consequential decisions shaped by a system they had no ability to see or challenge. The case became one of the most cited examples of what happens when algorithmic systems are adopted without adequate scrutiny of what those systems are actually measuring.

4. Workday and Algorithmic Hiring Bias

In 2023, a lawsuit was filed against Workday, one of the leading providers of enterprise HR software, alleging that its AI-powered screening tools discriminated against job applicants on the basis of age, race, and disability. The plaintiff, a job seeker who had been rejected by multiple employers using Workday's platform, argued that the screening algorithms had systematically disadvantaged him and others in similar situations.

The Workday case is notable because it raised a question that many organisations have been slow to confront: when an AI tool produces discriminatory outcomes, who is responsible? The employer who used the tool? The vendor who built it? The case was still working its way through the courts at the time of this writing, but the questions it raised have prompted a growing number of HR professionals to look more critically at the AI tools their organisations use for recruitment and performance management. Acemoglu and Restrepo (2018) had already warned that automation in labour markets could produce distributional consequences that policymakers and organisations were not ready to address. The Workday case was an instance of those consequences appearing in a legal rather than purely economic form.

WHEN AI COMPLICATES RATHER THAN SIMPLIFIES

Not all the complications that AI tools create rise to the level of a high-profile scandal. Most of them are quieter, experienced by individual employees in the course of ordinary working days, and rarely documented in ways that reach public attention. But they are nonetheless real and consequential for the organisations involved.

One common pattern is what might be called the verification trap. An employee who uses an AI tool to draft a report, generate a forecast, or summarise a dataset still needs to verify that the output is correct. If the tool is producing plausible-sounding text, the verification task can be harder than it would have been if the employee had simply done the work themselves. With their own work, they know where to look for errors. With an AI's output, the errors can be anywhere and can look entirely convincing. The cognitive cost of verification can easily exceed the cognitive cost of original production, particularly for tasks that require professional judgment or domain expertise.

A second pattern involves what Autor (2015) described as the hollowing out of intermediate-skill work. AI tools are generally most capable at the structured, rule-following aspects of a job. They are least capable at the aspects that require tacit knowledge, contextual judgment, or genuine creativity. This means that AI adoption can, in some roles, gradually shift the human portion of the work toward the most cognitively demanding tasks while simultaneously reducing the opportunity for employees to build the competence they need to perform those demanding tasks well. An employee who has been using AI to draft correspondence for two years may find that their own unaided writing has deteriorated, not because the AI was bad, but because the skill was not being exercised.

A third complication arises from what Wilson and Daugherty (2018) called the collaboration gap: the difficulty of integrating AI tool outputs smoothly into team workflows. When one employee uses an AI

tool to produce work that another employee needs to review, evaluate, or build on, questions quickly arise about provenance, accountability, and quality standards. Did the AI produce this, or did the person? How much was edited? If there is an error, whose error is it? These questions do not have easy answers, and teams that have not established clear norms around AI use often find themselves in unnecessary conflict or confusion.

Kahneman (2011), writing before the current wave of generative AI, distinguished between two modes of thinking: a fast, intuitive mode that operates on pattern recognition, and a slow, deliberate mode that is applied to problems that genuinely require careful reasoning. One of the risks of AI tool adoption is that tools which are designed to feel fast and effortless can encourage employees to apply fast, intuitive thinking to situations that actually require the slower, more deliberate kind. When an AI tool produces a confident-sounding answer in seconds, the psychological pull toward accepting that answer and moving on is real and difficult to resist, even when the stakes are high enough to warrant more careful thought.

ADVISORIES FOR WORKPLACE MANAGERS

The patterns described in the preceding sections are not inevitable. They arise from specific choices made by the organisations and managers responsible for AI tool adoption. The following advisories are addressed to those managers, and they are grounded in what the literature and the documented cases suggest actually makes a difference.

Advisory 1: Stop Treating AI Adoption as a Technical Project

Most AI tool deployments fail not because the technology is defective but because the human dimensions of adoption are treated as secondary. Managers who leave the rollout entirely to IT and procurement teams, and who assume that employees will adapt naturally once the tool is installed, consistently produce worse outcomes than managers who treat adoption as a change management challenge from the beginning. The technology acceptance research is unambiguous on this point. Perceived usefulness and perceived ease of use are not objective properties of a tool; they are perceptions formed by the employees who use it, shaped by training, by context, and by the quality of the explanation they receive about why the tool is being introduced and what it is supposed to do.

Advisory 2: Be Honest About What the Tool Cannot Do

One of the most reliably trust-damaging things an organisation can do when introducing an AI tool is oversell it. When employees encounter the tool's limitations in practice, after having been told it is transformative, the gap between promise and experience becomes the lens through which everything about the tool is subsequently interpreted. Managers should insist on honest, specific descriptions of the tool's capabilities and its known failure modes before any rollout communication is prepared. PricewaterhouseCoopers (2017) noted in its analysis of AI adoption that organisations which set realistic expectations in the early stages of AI deployment achieved higher sustained adoption rates than those that led with aspirational framing.

Advisory 3: Involve Employees in Evaluation Before Commitment

The employees who will use an AI tool are frequently in the best position to assess whether it fits their actual work. Managers who run genuine pilot evaluations, with real employees doing real tasks, and who listen seriously to what those employees report, are more likely to catch problems before they become costly. Equally important, employees who have been involved in the evaluation process approach the tool with a sense of informed ownership rather than imposed compliance. The research on technology adoption consistently shows that participation in the evaluation process predicts adoption quality, not just adoption rate.

Advisory 4: Create Visible Channels for Employees to Report AI Errors

The algorithm aversion effect is partly a product of isolation: an employee who witnesses an AI error and has nowhere to report it, no process for having the error investigated, and no sense that the organisation takes such errors seriously, will draw their own conclusions privately. Those conclusions tend to be more damning than the evidence warrants. Organisations that create clear, low-friction channels for employees to flag AI errors, and that demonstrably act on what they learn, accomplish two things at once. They get better information about tool performance, and they give employees a sense of agency over a process that can otherwise feel entirely out of their control.

Advisory 5: Build Skill Maintenance into Any AI-Assisted Workflow

When AI tools take over portions of a task that employees previously performed themselves, the skills involved in that task will atrophy over time if no one actively maintains them. Managers should identify which professional competencies are most at risk in any given AI deployment and build deliberate practice of those competencies into team workflows. This is not a Luddite argument against AI tools. It is a recognition that a professional who has lost the skills that allow them to critically evaluate an AI's output is no longer in control of their own work. Wilson and Daugherty (2018) argued persuasively that the most effective human-AI partnerships are those in which humans retain genuine expertise, because only genuine experts can identify when an AI system is operating outside its competence.

Advisory 6: Establish Accountability Before Something Goes Wrong

The Workday and COMPAS cases make the same point in different contexts: when an AI tool produces a harmful or discriminatory outcome, the absence of a clear accountability framework makes the damage worse. Managers who have not established in advance who is responsible for AI-assisted decisions, how those decisions can be reviewed and challenged, and what recourse is available to people adversely affected by them, will find themselves in a much more difficult position when something goes wrong. And something will eventually go wrong. This is not pessimism; it is what the error rate of any probabilistic system guarantees.

PRACTICAL IMPLEMENTATION OF THE ADVISORIES

Advisories are only useful if they can be translated into actions that managers can actually take. The following guidance addresses each advisory in practical terms.

Advisory 1: Making the Shift from Technical to Human-Centred Adoption

In practice, this means assigning ownership of the adoption process to someone whose primary expertise is in people and change management, not in technology. It means running structured conversations with employee groups before any tool is introduced, to understand what their current workflow looks like, where the pain points are, and what a genuinely helpful tool would need to do. It means designing onboarding that is specific to the actual tasks employees perform, rather than generic vendor-supplied training. And it means establishing a 90-day review point at which adoption quality is assessed, not just adoption rate.

Advisory 2: Operationalising Honest Expectation-Setting

Honest expectation-setting does not happen on its own. It requires a manager to push back on vendor claims, to request documented evidence of the tool's accuracy in contexts comparable to the organisation's own, and to ask specifically what the tool's known failure modes are and how frequently they occur. That information should then be communicated to employees in plain language before the tool is introduced. A simple one-page reference document, written from the employee's perspective and describing both what

the tool does well and where it should not be relied upon without checking, can do more for adoption quality than any amount of promotional communication.

Advisory 3: Running Meaningful Pilot Evaluations

A meaningful pilot evaluation is not a brief demonstration run by a vendor representative. It is a structured exercise in which a representative group of employees uses the tool for real tasks over a period of several weeks, keeps a simple record of their experiences, and participates in a facilitated debrief at the end. The debrief should ask specific questions: Where did the tool save time? Where did it create extra work? Where were outputs unreliable? What would need to change for the tool to be genuinely useful? The answers to these questions should directly influence the decision about whether to proceed, and if so, how.

Advisory 4: Building Functional Error Reporting

An error reporting channel that employees actually use needs to be simple, fast, and visibly connected to a response. A dedicated email inbox that nobody monitors is not a channel; it is a gesture. What works is a designated contact person who acknowledges reports within a defined timeframe, a regular review cycle in which reported errors are examined and categorised, and a feedback loop that tells the reporting employee what was found and what, if anything, will change. This does not require a large investment. It requires a clear assignment of responsibility and a commitment to taking reports seriously.

Advisory 5: Protecting Skills in Practice

Skill maintenance can be built into workflows in several practical ways. Teams that use AI tools for drafting can designate a proportion of their output, perhaps one in five documents, as human-only, with no AI assistance. Teams that use AI for data analysis can require analysts to produce their own independent assessment before consulting the AI's output, then reconcile the two. These practices need not be permanent; they are most important in the first year or two of AI adoption, when the risk of rapid skill atrophy is greatest. Managers should also watch for early warning signs: employees who are no longer able to perform a task without the AI tool, or who cannot explain why an AI output is or is not correct.

Advisory 6: Establishing Accountability Frameworks

An accountability framework for AI-assisted decisions does not need to be legally complex. At a minimum, it should specify who makes the final decision in any context where AI is providing input, what documentation is created to record that an AI tool was consulted and what it recommended, what the process is for a person who disputes a decision that was influenced by an AI tool, and what the escalation path looks like when an AI error causes harm. Many organisations have these frameworks for human decision-making but have not extended them to AI-influenced decisions. Closing that gap is a matter of policy design, not technology, and managers should not wait for a legal challenge to begin the work.

CONCLUSION

The question posed in this paper's title does not have a single answer. AI tools help some employees with some tasks in some contexts, and they complicate things for other employees in other contexts. What the literature and the documented cases together make clear is that the difference between those two outcomes is not primarily a function of the technology. It is a function of how the technology is introduced, communicated, supported, and governed.

Employees judge AI tool effectiveness through a practical, experiential lens. They ask whether a tool saves them time, whether its outputs are reliable, whether they can understand why it is telling them what it is telling them, and whether using it makes their working day feel easier or harder. When the answer to

enough of these questions is unsatisfying, trust erodes. Once trust has eroded, it is very difficult to restore, as the Watson Health and COMPAS cases both demonstrate.

The advisories offered in this paper are not exhaustive, and they are not painless to implement. They require managers to invest time, attention, and sometimes organisational capital in aspects of AI adoption that are easier to skip. But the cost of skipping them, measured in wasted investment, disengaged employees, and the kind of high-profile failures that damage organisations for years, is invariably higher than the cost of doing the work properly from the start.

AI is not going away, and neither is the responsibility of managers to make its adoption serve the people doing the actual work. That responsibility is more consequential than it might appear, and it deserves to be treated that way.

REFERENCES

- Acemoglu, D., & Restrepo, P. (2018). Artificial intelligence, automation, and work (NBER Working Paper No. 24196). National Bureau of Economic Research. <https://doi.org/10.3386/w24196>
- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3-30. <https://doi.org/10.1257/jep.29.3.3>
- Brynjolfsson, E., & McAfee, A. (2014). The second machine age: Work, progress, and prosperity in a time of brilliant technologies. W. W. Norton & Company.
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108-116.
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319-340. <https://doi.org/10.2307/249008>
- Dietvorst, B. J., Logg, J. M., & Soll, J. B. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114-126. <https://doi.org/10.1037/xge0000033>
- Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50-80. <https://doi.org/10.1518/hfes.46.1.50.30392>
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90-103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3), 709-734. <https://doi.org/10.2307/258792>
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230-253. <https://doi.org/10.1518/001872097778543886>
- PricewaterhouseCoopers. (2017). Sizing the prize: What's the real value of AI for your business and how can you capitalise? PwC. <https://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425-478. <https://doi.org/10.2307/30036540>
- Wilson, H. J., & Daugherty, P. R. (2018). Collaborative intelligence: Humans and AI are joining forces. *Harvard Business Review*, 96(4), 114-123.