The Leadership Quarterly 37 (2026) 101975

Contents lists available at ScienceDirect

The Leadership Quarterly

journal homepage: www.elsevier.com/locate/leaqua

Commentary

The renaissance of management: When regular employees become multi-level managers of AI agents

Niels Van Quaquebeke [a,b,c,*], Brooke A. Gazdag [b], Allan Lee [c]

[a] Kühne Logistics University, University of Exeter Business School, UK
[b] Kühne Logistics University, Germany
[c] University of Exeter Business School, UK

ARTICLE INFO

Keywords
Artificial intelligence
AI agents
Leadership
Management
Future of work
Digital

ABSTRACT

With Artificial Intelligence (AI) entering the white-collar workforce at scale, employees will increasingly operate multiple AI agents which, once aligned, can (semi-) autonomously handle a wide array of tasks. Our commentary explores this prospective reality and its implications for employees. Specifically, we argue that employees in these scenarios will effectively have to become vertical multi-level managers, that is, continuously code-switching between the top, middle, and lower management roles when interacting with their AI agents. Against this background, we outline anticipated benefits and potential challenges for the involved employees. Finally, we discuss the implications of this transformative shift for future management research and the requirements we foresee for educational practices within our field.

A new kind of organization is emerging—what some have called "frontier firms"—where employees increasingly operate multiple AI agents to get their work done (Altchek, 2025; Benioff, 2024; New York Times, 2025; Xu et al., 2025). Rather than completing tasks themselves, these employees coordinate a portfolio of AI systems, each capable of (semi-) autonomously executing goal-directed functions with high speed, precision, and uptime. While some companies still struggle to get simple large language models (LLMs) to support their workforce, following William Gibson's aphorism that "the future is already here—it's just not evenly distributed," real-world applications of effective AI agent integration are increasingly emerging also in traditional organizations (for an overview, see Renner & Chaban, 2025). The appeal of such setups is clear: managing increasingly agentic AIs enables greater efficiency and performance, resulting in higher profit (for a selection of support ranging from LLM to fully agentic AI see the following: Baek, Fu, Ma, & Peng, 2026; Brynjolfsson, Li, & Raymond, 2025; Cui et al., 2026; Suleyman & Bhaskar, 2023; Wang, Zhu, Feng, Lu, & Jia, 2026).

But AI agents are not just tools; they are (semi-) autonomous systems. As such, like human workers, they must be steered, aligned, and monitored. In short, they require *management* (Cohen, 2025). This presents a future our commentary seeks to explore in more depth.

Following Johns's (2024) and Oc's (2018) calls for more explicit contextual theorizing, we treat the rise of AI agents as a major contextual shift that calls for an evolution of existing management thinking. In response, and moving beyond the classic view in which management responsibilities are hierarchically *differentiated* across individuals, we argue that under these new forms of organizing work, employees will need to become vertically *integrated* multi-level managers of their AI agents—concurrently handling top-level strategic goals, middle-level coordination, and lower-level task oversight and prompting.

Of course, not all organizations or professions will experience these changes to the same extent, or on the same timeline, if at all (Acemoglu, Autor, Hazell, & Restrepo, 2022; Busse, Kach, & Wagner, 2016; Massenkoff & McCrory, 2026). The purpose of our commentary, therefore, is not to offer a general theory of organizations, but to engage in prospective, generative theorizing (Gümüsay & Reinecke, 2022) about one possible future of work as we can envision it from our vantage point in summer 2026. Our aim is to spark a forward-looking conversation about how leadership theory—and the way we educate future leaders—must evolve in anticipation of these new organizational forms, where managing AI agents may become a routine, even defining, feature of work.

* Corresponding author.
E-mail address: niels.quaquebeke@the-klu.org (N. Van Quaquebeke).

https://doi.org/10.1016/j.leaqua.2026.101975
Received 14 July 2025; Received in revised form 23 May 2026; Accepted 30 May 2026
Available online 10 June 2026
1048-9843/© 2026 Elsevier Inc. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

What Are AI Agents?

To understand the impending transformation, we first need to clarify what AI agents are. While many are already familiar with generative AI tools such as ChatGPT or Claude chat, AI agents represent a more advanced application, often encompassing multi-step planning and reasoning, as well as autonomous decision-making (Ferrag, Tihanyi, & Debbah, 2025). Unlike chatbots or passive assistants, AI agents are systems capable of acting toward goals with minimal human intervention. While definitions vary today and remain somewhat imprecise, depending on their type, AI agents can plan, have different decision authority, interact with multiple systems (e.g., via mimicking human behavior in interfaces or via APIs), learn, and adapt their behavior based on feedback from their environment (for an overview on different types of agents, see Gassmann & Wincent, 2025; for an overview of practical applications see Renner & Chaban, 2025). For example, an AI agent may be tasked with drafting a market analysis. To fulfill this task, it might search daily for relevant data, summarize findings, generate visualizations, and return a formatted report including recommendations for advertising spend or product adjustments. These capabilities make AI agents particularly suited to complex workflows that benefit from speed, consistency, and repeatability. At the same time, AI agents can also be designed to provide unsolicited input based on new information or to amplify weak signals that might otherwise remain undetected. In practice, employees are therefore likely to deploy multiple specialized AI agents—each focused on distinct tasks—operating in parallel or in layered, interdependent systems.

While still in early deployment, AI agents are advancing rapidly (Ferrag et al., 2025; Gassmann & Wincent, 2025; Renner & Chaban, 2025). As these systems grow more robust, the central question will shift from whether to use them to how best to manage them. The question of how humans manage these digital workers is therefore what we seek to explore in the present commentary.

Notably, while we speak of AI agents in general terms, we recognize that not all AI systems are equally "agentic." Our arguments apply most directly to systems that exhibit some degree of task autonomy, learning capacity, and interactive interfaces that enable ongoing oversight, coordination, or refinement. They apply less to narrowly scoped AI tools that are static, deterministic, or largely hidden behind the interface layer.

Why AI agents need to be managed, not led

As noted above, depending on their type, AI agents exhibit goal-oriented behavior underpinned by some degree of autonomy, adapt through continuous learning, and may generate probabilistic outputs, meaning they do not necessarily behave identically across tasks or iterations (e.g., customer agents, creative agents, research agents, etc.). This implies that the class of AI agents we focus on, despite being tools, require oversight that mirrors key dimensions of managing human workers, including thoughtful deployment, structured onboarding, clear instruction, boundary setting, ongoing monitoring, and adjustment based on feedback. They must be coordinated within workflows, aligned with broader objectives, and, at times, phased out or retrained. In short, what AI agents require is not a simple user—but a manager.

Notably, we deliberately frame the emerging challenge as one of management rather than leadership. While these terms are often used interchangeably in organizational life (see for an overview, Kniffin, Detert, & Leroy, 2020), the distinction matters in the context of human–AI collaboration. AI agents do not require inspiration, emotional resonance, or moral legitimacy, which are hallmarks of contemporary leadership theories. Nor do they respond to social influence, shared identity appeals, or charismatic vision.[1]

A short disclaimer on boundary conditions

Even though AI agents are projected to generally grow in capability, autonomy, and ubiquity (Stanford University Human-Centered Artificial Intelligence, 2025), the agentic AI transformation will certainly not unfold uniformly across organizations, industries, or geographies (Kergroach & Héritier, 2025). Adoption trajectories will vary depending on factors such as the nature of the work (e.g., regulatory constraints, need for embodied presence, or tacit judgment), market competition, sector-specific technological maturity, and differences in organizational size, culture, and resource availability. Likewise, individual differences will continue to play a role. Openness to change among executive boards and lower-level managers alike—along with digital literacy and risk tolerance—will all shape how readily multi-level AI management is embraced or resisted. National context and labor market pressures may further accelerate or inhibit adoption. In short, the shift we describe will not happen everywhere—or at the same pace.

Managing the machine workforce: A multi-level management model

In traditional organizations, strategic, tactical, and operational responsibilities are typically *distributed* across hierarchical layers of top, middle, and lower-level management. In contrast, as we will argue, employees operating AI agents will have to *integrate* all three levels of management within themselves. Notably, this type of proposed *vertical* integration of management roles is different from the combination of multiple *horizontal* management roles that can, for instance, be observed when traditional managers may oversee multiple distinct functional teams.

Of course, some employees have always needed to balance some form of vertical management, i.e., localized strategic and operational thinking. However, the rise of AI agents introduces qualitatively different coordination challenges. Human collaborators interpret intent, read emotional tone, and adapt flexibly to social and contextual cues. They negotiate, push back, or compensate for ambiguity. AI agents, in contrast, operate largely on the explicit information provided. Consequently, effective management of AI requires a level of structure, precision, and continuous oversight that would be counterproductive in human-only settings. Indeed, classical theories of management emphasize empowerment and avoidance of micromanagement to preserve motivation; yet AI agents, lacking motivation altogether, demand a clarity that may border on "micromanagement" to prevent error, drift, or ethical lapses.

Below, we describe what the proposed vertical multi-level management may look like in practice (see Table 1 for overview) before we

[1] Indeed, even the once fabled notion of "prompt engineering," where users were encouraged to tell LLMs, for example, that "their life depends on this answer" or user other emotional influence tactics (Li et al., 2023), are already becoming obsolete as AI systems learn to auto-prompt, that is, to automatically create, refine, and optimize prompts based on the objective they infer, or to ask clarifying questions to better understand that objective (e.g., Madaan et al., 2023; Meincke, Mollick, Mollick, & Shapiro, 2025; Zhuo et al., 2024). We would thus caution when trying to draw similarities between prompting and traditional people leadership communication.

Table 1
Overview of Roles for Multi-Level Managers when Supervising AI-Agent Portfolios.

Management Level	Core Purpose	Tasks & Responsibilities	Capabilities Needed	Cadence
Top-Level (Strategic Architect)	Steer the broader direction, principles, and guardrails of the entire agent ecosystem.	• Set long-term goals for which tasks or capabilities to automate. • Define the "why, when, and how much" of automation. • Establish and calibrate (ethical) guardrails and governance principles. • Decide on overall agent architecture and its expansion vs. reduction.	• Strategic foresight & systems thinking. • Ethical and responsible-AI judgment. • Risk-appetite calibration. • Goal/vision-setting	Quarterly to monthly direction-setting, with ad-hoc calibrations as environment or agent performance changes.
Middle-Level (Translator & Orchestrator)	Translate strategy into coherent systems and workflows; integrate diverse agents.	• Craft system prompts, chain tasks, and specify data flows. • Orchestrate handoffs between specialized agents. • Monitor cross-agent coherence and resolve conflicts. • Iterate workflow logic as context or data shifts.	• Systems thinking & process design. • Adaptive coordination & system conflict resolution. • Data-literate communication (prompt engineering). • Real-time system diagnostic reasoning.	Daily to hourly workflow tuning; coordination loops tighten whenever live data or agent outputs change.
Lower-Level (Operational Supervisor)	Guarantee quality, safety, and compliance at the output edge.	• Validate real-time outputs for accuracy and policy fit. • Detect anomalies or hallucinations; intervene instantly. • Maintain or fine-tune individual agents and curate feedback data. • Serve as final gatekeeper before external release.	• Acute attention to detail & pattern recognition. • Domain expertise for rapid plausibility checks. • Stress-resilient decision-making. • Accountability for compliance, integrity, and improvements.	Minute-to-minute monitoring during active runs; post-run reviews feed continuous learning loops.

Table note: Cadence is not only a matter of choice. Indeed, attending to a certain level may sometimes be reactive (i.e. forced upon the person by the AI system) or voluntary and self-initiated (i.e. a deliberate choice to attend to another task). There is a certain gravitation toward those levels that demand attention (i.e. lower level) in contrast to those that would need deliberate self-initiated self-regulation (top level) (Koch et al., 2018).

discuss the respective tension of benefits (i.e., gains) and challenges (i.e., pains). Note that while we follow the nomenclature of top, middle, and lower management, these do not denote hierarchical levels within our framework but only functional task categories rooted in the extant literature.

Top-level management of AI agents (strategic)

The highest level concerns the entire AI agent ecosystem, much like top executives or a supervisory board in a traditional organization. The human operator must define strategic objectives, including setting long-term goals for how agents are deployed, determining which capabilities to automate, and clarifying what might be called 'organizational design' to achieve these aims. In addition, the manager must establish ethical guidelines and governance principles to ensure transparency, fairness, and responsible use, especially as agents become more autonomous and gain greater decision latitude.

In this new model, even frontline employees managing AI agents will need to engage in this kind of strategic reflection. They may be consulted by AI systems on how best to manage their domain and will need to develop high-level deployment plans for their agents. For example, a product manager overseeing several AI agents might delegate competitor monitoring, market trend analysis, and customer sentiment tracking to specialized models, thereby freeing up cognitive bandwidth for more interpretive tasks such as feature prioritization. However, important decisions remain. Should an agent's recommendation on pricing be trusted outright, or should it be weighed against human judgment? Should a customer analysis agent surface every signal, or only anomalies based on preset thresholds?

Strategic oversight also includes ethical calibration. If one agent uses customer data to generate personalization insights, the employee must judge whether its data handling and output logic align with both internal policy and public expectations. For instance, if a model begins inferring sensitive demographic traits to increase engagement, the employee must decide whether that approach, even if technically feasible and effective, is ethically acceptable.

In traditional organizations, such governance questions might be escalated to legal or compliance teams. In agent-based work, however, these decisions are more likely to occur in real time and at the operational level, relying on the judgment of the individual employee. In this way, even frontline employees may find themselves in roles that resemble those of executives or supervisory boards, deciding not only how work is done, but also what kinds of work should be pursued.

This is not to suggest that ethical and legal rules will not exist, whether through organizational policies, contextual guardrails, or company-wide AI systems monitoring processes and outputs. Rather, as in contemporary organizations, rules and frameworks provide boundaries, but individuals must still exercise judgment in how they operate within them, particularly in contexts that demand speed and flexibility.

Middle-level management of AI Agents (translational and coordinative)

At the middle level, much like in traditional firms, the employee functions as a translator and integrator, bridging strategic intent and operational execution. While the high-level vision defines what is to be achieved, it must still be translated into how that vision is enacted by a diverse set of AI agents, each with its own strengths, limitations, and underlying logic.

Concretely, this may involve crafting base prompts and configuring task chains that convert abstract goals into actionable objectives. For instance, if the strategic goal is to improve customer retention, the employee may instruct one agent to segment customer data, generate predictive churn scores, and draft targeted outreach messages. At the same time, another agent may be tasked with analyzing the product or service itself, using past feedback to identify improvements and suggest ways to generate deeper customer insights. At this stage, the outputs of

both agents must be integrated, requiring them to effectively "interact" so that insights from product analysis can inform targeted outreach. Each of these steps requires not only technical specification, but also an understanding of workflow dependencies, including how information flows across agents, in what sequence, and with what handoffs.

To date, most agents are modular and domain-specific, and do not naturally collaborate. The employee must therefore act as an orchestration layer, ensuring that one agent's output is passed at the right moment and in the appropriate format to another. This involves designing not only isolated tasks, but also interconnected process logic, for example, through platforms such as n8n, AutoGen, or Kore.ai. In this sense, the role resembles that of a conductor coordinating different sections of a digital orchestra while monitoring their interplay.

This role also demands foresight. When coordinating multiple agents, the employee must anticipate potential failure points and conflicts between systems. What happens if the forecasting agent's predictions contradict the pricing agent's thresholds? Which agent takes priority? Which outputs are surfaced, and which are revised? Is the issue confined to a single agent, requiring lower-level intervention, or does it reflect a broader systemic problem that requires reconsideration of the overall logic? In short, once the AI agent system is established, the employee manages by exception and prioritizes iteration, intervening not to micromanage every detail, but to ensure system coherence and progressively improve overall performance.

Lower-level management of AI agents (operational and supervisory)

At the lower level, the employee serves as a real-time operational supervisor of individual agent behavior, ensuring accuracy, reliability, and safe execution at the output edge of the system. While middle-level management coordinates workflows between agents and aligns outputs with broader goals, operational management focuses on direct, ground-level oversight of each agent's execution. This involves monitoring live agent performance, identifying anomalies in output quality, and intervening when an agent goes off script. In contrast to the middle-level role, the task here is to determine whether an issue stems from the broader system logic, which may require strategic intervention, or from an individual agent that requires hands-on adjustment. For instance, if a content-generation agent begins producing with an inconsistent tone, or a data-cleaning agent fails to detect edge-case errors, the employee must quickly diagnose and correct the issue, whether by refining prompts, resetting the agent's state, or switching to a different model.

This role also includes the onboarding and ongoing maintenance of agents. When a new agent is deployed, for example, one trained to evaluate supplier proposals, the employee must verify that its logic aligns with existing criteria, that its recommendations are explainable, and that it integrates safely into the decision chain. Unlike middle-level integration, which focuses on sequencing and alignment between agents, this level of management ensures fitness for purpose at the point of task execution. Because agents evolve with use, the employee also plays a continuous feedback-loop role, adjusting what constitutes acceptable versus unacceptable outputs and reinforcing patterns accordingly. This may involve flagging outputs for retraining, curating feedback data, or disabling models that drift beyond acceptable performance thresholds. While AI systems may learn, it is the human who determines what counts as meaningful improvement.

Finally, in this operational role, employees serve as the last line of accountability before outputs leave the system, whether for internal or external use. This level of oversight is critical for safeguarding organizational performance and legal integrity.

Why all three management levels are needed within one employee

In this new paradigm, employees must be able to dynamically code-switch between all three layers of management. Some of this may be self-initiated, some of it simply a reaction to system demands. And each comes with its own switching costs. While all three management roles may sometimes be in demand simultaneously, we think they are more likely to unfold at different rhythms (see the cadence column in Table 1; for cognitive underpinnings, see Koch, Poljac, Müller, & Kiesel, 2018). Lower-level management tasks, such as monitoring execution, typically are very reactive and therefore often operate on a daily cadence. Moreover, because they are more reactive in nature, there is a natural gravitation toward and subsequent lingering attention around these tasks (cf. "biased competition" in Koch et al., 2018). In contrast, middle-level management tasks, such as delegating responsibilities and designing workflows, usually follow a more infrequent interval. And finally, tasks that fall under top-level management such as, setting long-term goals, guardrails, and principles, are often the result of infrequent but deliberate thought processes. Notably, operating on this level of abstraction is not prompted and thus requires more self-regulation, i.e., self-initiation.

This integration of management levels within a single dynamic role marks a fundamental departure from traditional organizational design of discrete waterfall hierarchies. It demands both vertical agility (switching between abstract goals and concrete tasks) and horizontal dexterity (managing interactions across a suite of different AI agents). In short, the employee becomes an all-in-one strategist, coordinator, and supervisor—responsible for the objectives, conduct, and coherence of their digital team: A multi-level manager.

The tensions of gains and pains in multi-level AI management

The consolidation of multiple layers of management within individual employees offers distinct advantages but also introduces structural and psychological challenges. Below, we examine key aspects of this shift, showing how each potential gain is accompanied by a corresponding cost, resulting in six tensions that current management theorizing and education have yet to fully address (see Table 2 for an overview). Notably, this list is not derived from a single organizing framework, nor is it intended to be exhaustive. Rather, through an abductive process within the author team and iterative refinement during the review process, it reflects the tensions we see as central to this emerging transformation. Broadly, these tensions move from the cognitive and social demands of the role, through its psychological costs, to deeper questions of agency, meaning, and ethical responsibility.

Streamlined control vs. cognitive overload

With one employee responsible for strategic framing, coordination, and execution, accountability becomes clearer. Decisions and outcomes can be traced to a single source, reducing ambiguity and improving ownership. As a consequence, workflows become more agile and decisions more coherent (Park, Baer, & Nickerson, 2025). In theory, this type of internal alignment may allow employees to not only produce more but also liberate their attention for higher-order thinking, creativity, and judgment.

Yet this total control comes at a cost: managing across levels simultaneously increases mental load. Constantly toggling between big-picture thinking and operational troubleshooting can create role confusion, decision fatigue, and psychological strain as the complexity grows (Steinbach, Gamache, & Johnson, 2019) and even more so when these switches are experienced as reactive rather than deliberately self-initiated (Koch et al., 2018). And while the ensuing execution speed may feel empowering, it also destabilizes. When work unfolds faster than humans can interpret or contextualize it, reflection collapses into reaction and oversight becomes superficial (Weick, 1988). As output accelerates, so do accountability demands. In consequence, employees find themselves supervising systems they can no longer fully track, let alone pause to interrogate (cf. "brain fry", Bedard et al., 2026). And as the organizational beat accelerates, so does employees' heart rate. Cognitive

Table 2
Selection of Six Tensions of Gains and Pains of Multi-Level AI Agent Management.

Dimension	Upside	Downside
Accountability	Streamlined control – ownership and traceability sit with one human orchestrator.	Cognitive overload – the always-on mode and the constant code-switching breeds decision fatigue and psychological strain.
Social Dynamics	Logical technocracy– logic instead of politics and charisma, leveling the playing field for quieter, competent voices.	Social atrophy – less practice in empathy, social learning, conflict resolution weakens people skills over time.
Productivity	Capacity – one person can leverage agents to significantly expand where and how deep they can meaningfully contribute.	Loss of meaning – personal accomplishment can erode when machines do most of the heavy lifting.
Dominion	Autonomy – access to skills and a command of knowledge outside one's domain allow for more independent work.	Misalignment – organizational demarcation lines of responsibilities blur as people's AI agents encroach on each other's domains, resulting in coordination and ownership issues.
Agency	Strategic empowerment – resources previously reserved for higher-level management (e.g., lower-level consultants, data, access to evidence) are readily available to inform decision-making.	Erosion of agency – human strategic judgment becomes increasingly passive and may atrophy, disempowering employees in the larger AI-agent ecosystem.
Ethics	Ethical consistency – a single decision-maker sets and enforces moral guardrails.	Ethical isolation – without peer debate, questionable choices may go unchallenged; less representation in perspectives.

strain turns into psychological stress (Ranganathan & Ye, 2026). Employees start to miss the comforting colleagues, their contrarian sparring partners who challenge and thereby productively slow down the process (Shaw et al., 2011), or the teammate who signals it is time to stop iterating and move forward (Gerpott, Lehmann-Willenbrock, Voelpel, & Van Vugt, 2019). Without other humans in the loop, the burden of control rests entirely on the shoulders of one individual. That may simply be too much to handle for some.

For example, a product manager takes pride in continually expanding their suite of AI agents. Total control over processes that once needed many frustrating alignment meetings. On paper, each deployed agent is aligned and should therefore free up capacity for more strategic thinking. Yet, while everything moves faster, their cognitive bandwidth does not. In reality, the torrent of drafts, simulations, and recommendations demands constant triage. But pausing to reflect begins to feel like falling behind. Even as their smartwatch signals rising stress levels and prompts them to take a moment to breathe, there seems to be no time to do so.

Logical technocracy vs. social atrophy

As outlined above, AI agents respond primarily to logic rather than emotion. When interacting with such systems, employees are less reliant on motivating, persuading, or managing interpersonal sensitivities. This can reduce the need for political gamesmanship and charismatic influence, and may level the playing field for less extroverted individuals who do not rely on interpersonal signaling to succeed.

However, over time, this relative shift toward technocratic engagement may contribute to a gradual erosion of social skills (cf. social atrophy, Kirk, Gabriel, Summerfield, Vidgen, & Hale, 2025; Van Quaquebeke & Gerpott, 2024). Managing machines does not train empathy, conflict resolution, or perspective-taking, and may even suppress these capabilities (Tey, Mazar, Tomaino, Duckworth, & Ungar, 2024). If such skills are not regularly exercised, they may deteriorate. Similarly, when knowledge is easily generated and readily available, opportunities for social learning may decline, weakening the foundations of interpersonal development (Denrell, Luukkonen, Chater, & Liu, 2026). As a result, returning to human interaction and coordination may feel increasingly effortful or inefficient and may ultimately be avoided (Hoffmann, Boysel, Nagle, Peng, & Xu, 2024).

For example, consider a scenario where an HR specialist oversees AI systems that screen résumés, conduct and score interviews, predict cultural fit, and monitor onboarding. Hiring bias drops (Ip, 2025) and efficiency rises, but human touch and contact fade. Over time, she grows uneasy handling difficult conversations with rejected candidates or mediating team conflicts—her interpersonal skills once honed through messy human interaction have slowly atrophied.

Capacity vs. loss of meaning

By providing instant access to specialized knowledge, analytical capabilities, and decision-making support, AI enables individuals to operate effectively across domains that would otherwise be beyond their expertise. This augmentation allows employees to tackle more complex, cross-functional challenges, contribute at higher levels of abstraction, and deliver value in areas they may not have been trained for. The result is not just faster work, but qualitatively richer and more versatile performance (Dell'Acqua et al., 2023)—expanding what it means to be a capable contributor.

Yet this expansion of capacity can diminish the sense of personal accomplishment (Kobiella, Flores López, Waltenberger, Draxler, & Schmidt, 2024). What is one's own contribution when it feels like the AI agents are doing the heavy lifting? When agents outperform humans, pride can give way to ambivalence or even anxiety (Kreft, Boguszewicz-Kreft, & Fydrych, 2023). If everything can be done faster and better by machines, what remains of human contribution? Why even strive when it is essentially not you delivering the performance?

For example, a consultant uses AI copilots to synthesize data, model scenarios, and generate client-ready slide decks overnight. Projects that used to demand late nights and teamwork now require only a few precise prompts. Billable output doubles, but hallway excitement vanishes. When clients praise “her” insights, she smiles politely—knowing most of the brilliance came from elsewhere.

Autonomy vs. misalignment

As noted above, AI agents blur traditional boundaries of expertise and responsibility. Employees equipped with AI agents can in principle perform tasks far outside their formal domain without needing others—designers prototype code, marketers generate data analyses, engineers draft user research reports. At first glance, this autonomy is deeply empowering. It overcomes bottlenecks and speeds up iteration cycles.

Yet, this newfound autonomy collides with established structures of coordination and accountability. If everyone can do everything, who is actually in charge of what? Without explicit governance templates or shared mental models, alignment devolves into ad-hoc negotiation. As such, personal autonomy may expand faster than organizational structures to contain it, creating role ambiguity within and role conflict between people, duplicated efforts, and latent tension between ownership and collaboration.

For example, in a consumer-tech company, a small product team uses AI agents for coding, design, and market testing. The designer now prototypes interfaces *and* edits backend code; the marketer runs A/B tests *and* rewrites UX copy. Everyone works with remarkable independence—until integration meetings reveal the trespasses and

territorial trenches are dug.

Strategic empowerment vs. erosion of agency

AI agents offer employees unprecedented access to strategic decision-making quality—delineating and triangulating decisions, evidence-based and data-driven. This unlocks a new form of strategic empowerment, once reserved for higher-level management, who could draw upon their whole organization (and outside strategy consultants) to improve their decision quality.

Yet with that empowerment comes a risk of strategic agency erosion. As AI systems propose ever more optimized solutions—instantly and confidently and, most importantly, increasingly even better than humans (Patwardhan et al., 2025)—human input may shrink to mere approval, rubber-stamping outputs without engaging in real reflection (Alon-Barkat & Busuioc, 2023) or independent knowledge generation (Denrell et al., 2026). Over time, the perceived necessity (and even capacity) for human strategic judgment may disappear, leaving employees functionally disempowered—even as they remain formally responsible.

For example, a government analyst relies on AI foresight tools to draft policy options based on decades of data and modeled outcomes. The system's analyses are so comprehensive and complex that human deliberation seems redundant or too cumbersome. Meetings become shorter; signoffs faster. Eventually, he realizes he has stopped questioning the recommendations altogether.

Ethical clarity vs. ethical isolation

With no need to persuade a team, ethical decisions rest squarely with the human AI manager. This allows for a stringency in the fairness or ethics logics that have a tendency to become inconsistent and provide ample ground for dissatisfaction when multiple parties with different fairness norms seek to be represented (Giessner & Van Quaquebeke, 2010).

Apart from the non-representation of certain ethical perspectives, this clarity can also create a form of ethical solitude for the individual employee. Without the social checks and discussions typical of human teams, questionable choices may go unchallenged or even praised (cf. sycophany, Cheng et al., 2026; Glickman & Sharot, 2024; Rathje et al., 2025). Ethics becomes less about negotiating consensus and more about internalized restraint (i.e. integrity)—a responsibility and accountability not all are prepared to carry alone.

For example, a news platform's editor supervises AI agents that filter and rank stories for millions of readers. The algorithm is programmed to enforce strict neutrality so that bias complaints drop to almost zero. Still, when a sensitive story about potential political corruption disappears from the feed, the editor must face the moral crisis all alone.

After this listing, we should note that we do not suggest that the tensions outlined here will manifest uniformly across all organizational contexts. As already alluded to in the above section on "a short disclaimer about boundary conditions", their salience will depend on contextual factors such as the degree of AI integration, organizational structure, task complexity, and the governance arrangements surrounding AI use. Consequently, in different organizations and roles, employees may experience some tensions stronger than others and sometimes not at all.

Implications for our field

Management theory and education need to acknowledge the emerging reality of a managerial hierarchy collapsed into a one-person, many-agent nexus. This demands a) new research to understand its implications and b) new pedagogical foci to equip individuals for these roles.

Implications for future leadership/management research.

In the future we outlined, researchers will need to reconceptualize AI agent management as less of a distributed and more of an integrated process with (semi-) autonomous, highly skilled but non-sentient collaborators. Several key questions follow.

What does it mean to manage?

Foundational management frameworks may need substantial revision. Take Mintzberg's classic typology, which grouped ten managerial roles into three categories: interpersonal, informational, and decisional. Traditionally, managers have spent up to 80% of their time on interpersonal communication (Mintzberg, 1973). Indeed, we have seen updates to what is required of managers when COVID hit and interactions went remote (Garretsen, Stoker, Soudis, & Wendt, 2024). But what happens when much of that interaction shifts from humans to machines? In such a world, informational and decisional roles may come to dominate daily managerial work, while interpersonal roles are minimized, redefined, or perhaps even rendered obsolete.

How are hierarchies and status negotiated?

When a single employee performs managerial tasks traditionally spread across multiple organizational layers, hierarchical distinctions blur (cf. self-managed organizations Lee & Edmondson, 2017). This raises new questions about how status, authority, and contribution are evaluated in AI-integrated firms (for a traditional account see Magee & Galinsky, 2008). What are the career prospects (i.e., promotions) for employees in such firms? Who is a "lower level" employee when effectively everyone is managing a fleet of AI agents—is it a simple counting game of how many agents one operates, or will we see hierarchical differentiation according to new responsibilities? How do employees perceive their own value when most of their output is mediated through machines? What new signals of competence or status emerge in a world driven by largely invisible orchestration and curation of machines?

What are the psychological consequences of a multi-level managerial role?

Multi-level management may also have complex psychological effects. On one hand, constantly cognitively shifting between strategic, tactical, and operational modes may produce role conflict, cognitive fatigue, and even some degree of identity ambivalence (Ashforth, Kreiner, & Fugate, 2000; Leroy, Schmidt, & Madjar, 2020; Monsell, 2003). On the other hand, it may lead to more integrated work identities, where employees feel greater autonomy, holistic control, and personal authorship over outcomes. This duality invites not only general questions about identity experiences but also more specifically whether managing AI agents could enhance cognitive flexibility and integrative thinking—or exhaust it.

How to leave an AI team?

Another unexplored dimension is the emerging bond between the employee and their AI stack. Over time, employees will train, refine, and customize their agents—embedding personal, relational, and institutional knowledge into them. But unlike a team that can follow a manager across firms, these agents often remain property of the organization. Will employees begin to feel locked in—not by loyalty or contracts, but by the compounded *digital capital* they cannot take with them? This not only provides a new human capital management perspective (Coff, 1997), but it will likely also bring about new psychological dynamics beyond established mobility constraints (Campbell, Coff, & Kryscynski, 2012) when high performers start to feel algorithmically handcuffed as they (due to their AI agent teams) find themselves unable to leave. Conversely, there are new questions to explore regarding knowledge management (Argote & Ingram, 2000) when a team of AI agents is to be handed over once its manager leaves, but potentially highly personalized traces can be found in each of the agents.

What will the different layers of collaboration in hybrid human–AI–human models look like?

The shift to AI agent management may require new forms of coordination and control between humans (Kellogg, Valentine, & Christin, 2020). When multiple employees each operate their own agent ecosystems, the human collaboration layer then becomes an interaction between humans nested in their web of digital actors. What kinds of misunderstandings, synchrony, or emergent structures will arise in such environments? Will psychological ownership over their finely tuned stack spark human employees to defend their digital teams much like coaches guard their athletes and —perhaps because they are their own creations—even more so than today's human managers defend their human teams (Ashforth & Lee, 1990)? And will an eventual atrophy of social skills amplify these tensions because people are not used to navigating human-human interaction anymore?

How to preserve human engagement in automated workflows?

As AI agents increasingly take over substantial portions of traditional workflows, an open question emerges regarding how human operators can maintain a meaningful sense of ownership, accountability, and agency over their work when much of the "doing" is abstracted away? What practices, rituals, or interface designs help reinforce a sense of control and contribution? These are, among others, central questions for work design research (Parker & Grote, 2022), decision-making control (Shrestha, Ben-Menahem, & Von Krogh, 2019), but also the human leadership of the employees (Raisch & Krakowski, 2021). Understanding how human managers of AI systems can remain experientially connected to their output is critical—not just for well-being, but also for motivating responsible oversight in increasingly automated environments.

What are the new units of analysis?

These shifts in practice also call for corresponding shifts in research methodology. As employees increasingly operate through AI agents, the traditional unit of analysis—the individual—may no longer suffice. Researchers will need to ask: how should we study teams when collaboration unfolds not only between people, but between people-plus-agent ecosystems (Orlikowski & Scott, 2008)? Likewise, how can the outlined multi-level managerial work be meaningfully captured when much of it is distributed across prompts, model outputs, and alignment orchestration? New frameworks will be required to not just assess system effectiveness but also system-level coherence, agent performance, and cross-agent integration, i.e. system efficiency, inferred from system error traces and token use.

Implications for management education and training

Preparing employees in AI-integrated organizations demands more than updated technical training—it calls for a reimagining of what it means to manage as a multi-level manager of AI agents.

Toward multi-level managerial literacy and agency

First, management education must move beyond siloed instruction. Instead, curricula should emphasize systemic thinking (Atwater, Kannan, & Stephens, 2008; Parker & Grote, 2022)—helping learners understand how top-level goal setting, middle-level coordination, and operational oversight interlock in real time. Managing AI agents requires fluency across these levels, and training must reflect that complexity. This means teaching not just what managers do, but also how each level feeds into each other, and, most importantly, how to seamlessly shift between hierarchy levels without losing coherence or control.

On this last point of control, we share the concerns raised by thinkers such as Yuval Harari and Sam Altman, who warn that one of AI's greatest risks is relegating humans to mere symbolic approval roles, where humans—whether out of a) convenience, b) lack of understanding, or c) risk aversion—inadvertently strip themselves of substantive strategic authority (Altman, & Bowman, M., 2025; Harari, 2017). Management education of the future needs to tackle all three if it wants to ensure that humans stay in the driver's seat (Raisch & Krakowski, 2021). As such, it seems key to encourage people not just to understand how a model arrived at its output (cf., "explainable AI", as currently mandated by various AI policy acts) but also to train them in critically engaging with the presented reasoning process itself, including spotting and challenging model assumptions and premises on every level of the managerial hierarchy (Larson, Moser, Caza, Muehlfeld, & Colombo, 2024).

Training for cognitive flexibility and adaptivity

Second, as prerequisite for the above, pedagogy must shift toward cultivating cognitive flexibility and adaptive learning (Patil, Srinivas, Tussing, & Rhee, 2025). Managing AI agents involves rapid perspective-switching, constant recalibration, and distributed attention across abstract strategy and granular execution. To prepare for this, learning environments should simulate these demands by providing high-fidelity experiences using AI-powered tools, live dashboards, and agent orchestration exercises. Such exercises then also need to cultivate a deeper focus on self-leadership and metacognitive regulation—the capacity to monitor one's own cognitive load across levels, recognize when strategic altitude or operational depth is required, and thus, deliberatively redeploy attention, for instance, to more abstract levels rather than letting oneself be constantly drawn to the concrete operational level. The goal is to help future employees navigate this new mental terrain without incurring excessive psychological switching costs.

Creating anchors for identity and meaning

Third, and perhaps most critically, organizations and educators must invest in helping employees reflect on and craft deeper meaning for their work and identity (Pratt & Ashforth, 2003). In AI agent environments, employees may no longer derive fulfillment from direct creation or team recognition, but will only have design, oversight, and curation to show for it. While such a shift can already be destabilizing for regular managers having been promoted out of their teams, the new reality will present even larger challenges. As such, reflective spaces should invite learners to explore what it means to lead without committed followers but only compliant machines, how to tune one's own compass for a true North when there are no other humans to challenge it, and how to cultivate a sense of belonging to a social system of machines.

Conclusion: Embracing the management renaissance

Much like how early automobiles were first described as "horseless carriages," there is naturally a risk of misrepresenting the transformation by viewing it through the lens of what came before (Banks, Dionne, Mast, & Sayama, 2022). Indeed, in the future, we may have to move beyond analogies that simply map the new terrain onto existing leadership or team models (Ramaul, Ritala, Kostis, & Aaltonen, 2025). For instance, with newer versions of agentic AI, where a user may only need to deploy a single agent that then creates and supervises a team of other agents (e.g., Claude Code Team Mode), these developments raise the question of what aspects of management will ultimately remain human (Van Quaquebeke & Gerpott, 2023). Admittedly, when it comes to such a yet unrealized future, our crystal ball in the year 2026 is no clearer than anyone else's.

For now, therefore, our analysis remains within the confines of Kniffin et al.' (2020) call to revalue management in its own right. Rather than signaling the end of management, the rise of advanced AI agents may instead mark the beginning of its renaissance—at least before the next epoch of AI is ushered in.

Author note

In keeping with the theme of our paper, we used ChatGPT to support

idea development and sharpen the prose, while we were mindful to preserve our own voice and lines of argumentation. Additionally, we used ChatGPT to provide reviewer feedback on the paper, which we incorporated where appropriate. This type of co-creation follows the spirit of recent calls to reimagine how research is conducted (Van Quaquebeke, Tonidandel, and Banks, 2025). Because the present paper concerns AI agents, we should add, for the avoidance of doubt, that no AI system made autonomous decisions regarding the paper at any point. The humans remained firmly in charge, for better or worse

CRediT authorship contribution statement

Niels Van Quaquebeke: Writing – review & editing, Writing – original draft, Conceptualization. **Brooke A. Gazdag:** Writing – review & editing, Conceptualization. **Allan Lee:** Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *J. Labor Econ., 40*(S1), S293–S340. https://doi.org/10.1086/718327

Alon-Barkat, S., & Busuioc, M. (2023). Human–AI interactions in public sector decision making: "Automation Bias" and "selective adherence" to algorithmic advice. *J. Public Adm. Res. Theory, 33*(1), 153–169. https://doi.org/10.1093/jopart/muac007

Altchek, A. (2025). Nvidia CEO says there's "no question" that we'll all be working alongside AI employees. *Bus. Insid.* https://www.businessinsider.com/nvidia-ceo-ai-employees-workers-2024-11.

Altman, S., & Bowman, M. (Directors).. (2025). *OpenAI CEO Sam Altman speaks with fed's Michelle Bowman on bank capital rules* [video recording]. YouTube https://www.youtube.com/watch?v=tScbQiavmpA.

Argote, L., & Ingram, P. (2000). Knowledge transfer: A basis for competitive advantage in firms. *Organ. Behav. Hum. Decis. Process., 82*(1), 150–169. https://doi.org/10.1006/obhd.2000.2893

Ashforth, B. E., Kreiner, G. E., & Fugate, M. (2000). All in a day's work: Boundaries and Micro role transitions. *Acad. Manag. Rev., 25*(3), 472. https://doi.org/10.2307/259305

Ashforth, B. E., & Lee, R. T. (1990). Defensive behavior in organizations: A preliminary model. *Hum. Relat., 43*(7), 621–648. https://doi.org/10.1177/001872679004300702

Atwater, J. B., Kannan, V. R., & Stephens, A. A. (2008). Cultivating systemic thinking in the next generation of business leaders. *Acad. Manag. Learn. Educ., 7*(1), 9–25. https://doi.org/10.5465/amle.2008.31413859

Baek, J., Fu, Y., Ma, W., & Peng, T. (2026). *AI agents for inventory control: Human-LLM-OR complementarity (version 2)*. arXiv. https://doi.org/10.48550/ARXIV.2602.12631.

Banks, G. C., Dionne, S. D., Mast, M. S., & Sayama, H. (2022). Leadership in the digital era: A review of who, what, when, where, and why. *Leadersh. Q., 33*(5), Article 101634. https://doi.org/10.1016/j.leaqua.2022.101634

Bedard, J., Kropp, M., Hsu, M., Karaman, O. T., Hawes, J., & Kellerman, G. R. (2026). *When Using AI Leads to "Brain Fry."* *Harvard Business Review*. https://hbr.org/2026/03/when-using-ai-leads-to-brain-fry.

Benioff, M. (2024). *How new digital workers will lead to an unlimited age*. TIME. https://time.com/7178872/agents-unlimited-age/.

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *Q. J. Econ., 140*(2), 889–942. https://doi.org/10.1093/qje/qjae044

Busse, C., Kach, A. P., & Wagner, S. M. (2016). Boundary conditions: What they are, how to explore them, why we need them, and when to consider them. *Organ. Res. Methods*. https://doi.org/10.1177/1094428116641191

Campbell, B. A., Coff, R., & Kryscynski, D. (2012). Rethinking sustained competitive advantage from human capital. *Acad. Manag. Rev., 37*(3), 376–395. https://doi.org/10.5465/amr.2010.0276

Cheng, M., Lee, C., Khadpe, P., Yu, S., Han, D., & Jurafsky, D. (2026). Sycophantic AI decreases prosocial intentions and promotes dependence. *Science, 391*(6792), eaec8352. https://doi.org/10.1126/science.aec8352

Coff, R. W. (1997). Human assets and management dilemmas: Coping with hazards on the road to resource-based theory. *Acad. Manag. Rev., 22*(2), 374. https://doi.org/10.2307/259327

Cohen, N. (2025). *The skills we used to teach managers? Now everyone needs them.* Forbes. https://www.forbes.com/sites/niritcohen/2025/05/17/the-skills-we-used-to-teach-managers-now-everyone-needs-them/.

Cui, K. Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., & Salz, T. (2026). The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers. *Manag. Sci.*. https://doi.org/10.1287/mnsc.2025.00535. mnsc.2025.00535.

Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., … Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *SSRN Electron. J.*. https://doi.org/10.2139/ssrn.4573321

Denrell, J., Luukkonen, J., Chater, N., & Liu, C. (2026). The need for absorptive capacity alleviates the free-rider problem in knowledge production. *Manag. Sci.*. https://doi.org/10.1287/mnsc.2024.06471. mnsc.2024.06471.

Ferrag, M. A., Tihanyi, N., & Debbah, M. (2025). *From LLM reasoning to autonomous AI agents: A comprehensive review (version 1)*. arXiv. https://doi.org/10.48550/ARXIV.2504.19678.

Garretsen, H., Stoker, J. I., Soudis, D., & Wendt, H. (2024). The pandemic that shocked managers across the world: The impact of the COVID-19 crisis on leadership behavior. *Leadersh. Q., 35*(5), Article 101630. https://doi.org/10.1016/j.leaqua.2022.101630

Gassmann, O., & Wincent, J. (2025). The non-human Enterprise: How AI agents reshape organizations. *California Management Review Insights.* https://cmr.berkeley.edu/2025/10/the-non-human-enterprise-how-ai-agents-reshape-organizations/.

Gerpott, F. H., Lehmann-Willenbrock, N., Voelpel, S. C., & Van Vugt, M. (2019). It's not just what is said, but when it's said: A temporal account of verbal behaviors and emergent leadership in self-managed teams. *Acad. Manag. J., 62*(3), 717–738. https://doi.org/10.5465/amj.2017.0149

Giessner, S. R., & Van Quaquebeke, N. (2010). Using a relational models perspective to understand normatively appropriate conduct in ethical leadership. *J. Bus. Ethics, 95*(S1), 43–55. https://doi.org/10.1007/s10551-011-0790-4

Glickman, M., & Sharot, T. (2024). How human–AI feedback loops alter human perceptual, emotional and social judgements. *Nat. Hum. Behav., 9*(2), 345–359. https://doi.org/10.1038/s41562-024-02077-2

Gümüsay, A. A., & Reinecke, J. (2022). Researching for desirable futures: From real utopias to imagining alternatives. *J. Manag. Stud., 59*(1), 236–242. https://doi.org/10.1111/joms.12709

Harari, Y. N. (2017). *Homo Deus: A brief history of tomorrow*. HarperCollins Publishers.

Hoffmann, M., Boysel, S., Nagle, F., Peng, S., & Xu, K. (2024). Generative AI and the nature of work. *SSRN*. https://doi.org/10.2139/ssrn.5007084

Ip, E. (2025). Fair AI in hiring: Experimental evidence on how biased hiring algorithms and different debiasing methods affect the quality and diversity of applicants. *Behav. Sci. Policy, 11*(1), 44–54. https://doi.org/10.1177/23794607251353585

Johns, G. (2024). The context deficit in leadership research. *Leadersh. Q., 35*(1), Article 101755. https://doi.org/10.1016/j.leaqua.2023.101755

Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Acad. Manag. Ann., 14*(1), 366–410. https://doi.org/10.5465/annals.2018.0174

Kergroach, S., & Héritier, J. (2025). *Emerging divides in the transition to artificial intelligence* (147th ed.). https://doi.org/10.1787/7376c776-en OECD regional development papers) [OECD regional development papers].

Kirk, H. R., Gabriel, I., Summerfield, C., Vidgen, B., & Hale, S. A. (2025). Why human–AI relationships need socioaffective alignment. *Humanities & Social Sciences Communications, 12*(1), 728. https://doi.org/10.1057/s41599-025-04532-5

Kniffin, K. M., Detert, J. R., & Leroy, H. L. (2020). On leading and managing: Synonyms or separate (and unequal)? *Academy of Management Discoveries, 6*(4), 544–571. https://doi.org/10.5465/amd.2018.0227

Kobiella, C., Flores López, Y. S., Waltenberger, F., Draxler, F., & Schmidt, A. (2024). "If the machine is as good as me, then what use am I?" – How the use of ChatGPT changes young professionals' perception of productivity and accomplishment. *Proceedings of the CHI Conference on Human Factors in Computing Systems, 1–16*. https://doi.org/10.1145/3613904.3641964

Koch, I., Poljac, E., Müller, H., & Kiesel, A. (2018). Cognitive structure, flexibility, and plasticity in human multitasking—An integrative review of dual-task and task-switching research. *Psychol. Bull., 144*(6), 557–583. https://doi.org/10.1037/bul0000144

Kreft, J., Boguszewicz-Kreft, M., & Fydrych, M. (2023). (lost) pride and prejudice. Journalistic identity negotiation versus the automation of content. *Journal. Pract.*, 1–24. https://doi.org/10.1080/17512786.2023.2289177

Larson, B. Z., Moser, C., Caza, A., Muehlfeld, K., & Colombo, L. A. (2024). Critical thinking in the age of generative AI. *Acad. Manag. Learn. Educ., 23*(3), 373–378. https://doi.org/10.5465/amle.2024.0338

Lee, M. Y., & Edmondson, A. C. (2017). Self-managing organizations: Exploring the limits of less-hierarchical organizing. *Res. Organ. Behav., 37*, 35–58. https://doi.org/10.1016/j.riob.2017.10.002

Leroy, S., Schmidt, A. M., & Madjar, N. (2020). Interruptions and task transitions: Understanding their characteristics, processes, and consequences. *Acad. Manag. Ann., 14*(2), 661–694. https://doi.org/10.5465/annals.2017.0146

Li, C., Wang, J., Zhang, Y., Zhu, K., Hou, W., Lian, J., Luo, F., Yang, Q., & Xie, X. (2023). *Large language models understand and can be enhanced by emotional stimuli (version 7)*. arXiv. https://doi.org/10.48550/ARXIV.2307.11760.

Madaan, A., Tandon, N., Gupta, P., Hallinan, S., Gao, L., Wiegreffe, S., … Clark, P. (2023). *Self-refine: Iterative refinement with self-feedback (version 2)*. arXiv. https://doi.org/10.48550/ARXIV.2303.17651.

Magee, J. C., & Galinsky, A. D. (2008). Social hierarchy: The self-reinforcing nature of power and status. *Acad. Manag. Ann., 2*(1), 351–398. https://doi.org/10.5465/19416520802211628

Massenkoff, M., & McCrory, P. (2026). *Labor market impacts of AI: A new measure and early evidence [Industry Report]*. Anthropic.

Meincke, L., Mollick, E., Mollick, L., & Shapiro, D. (2025). *Prompting science report 2: The decreasing value of chain of thought in prompting (version 1)*. arXiv. https://doi.org/10.48550/ARXIV.2506.07142.

Mintzberg, H. (1973). *The nature of managerial work*. Harper & Row.

Monsell, S. (2003). Task switching. *Trends Cogn. Sci., 7*(3), 134–140. https://doi.org/10.1016/S1364-6613(03)00028-7

New York Times. (2025). *A.I. Will transform the global economy*. The New York Times. https://www.nytimes.com/2024/12/07/special-series/ai-transform-global-economy.html.

Oc, B. (2018). Contextual leadership: A systematic review of how contextual factors shape leadership and its outcomes. *Leadersh. Q., 29*(1), 218–235. https://doi.org/10.1016/j.leaqua.2017.12.004

Orlikowski, W. J., & Scott, S. V. (2008). Sociomateriality: Challenging the separation of technology, work and organization. *Acad. Manag. Ann., 2*(1), 433–474. https://doi.org/10.5465/19416520802211644

Park, C. H., Baer, M., & Nickerson, J. (2025). Looking at the trees to see the Forest: Construal level shift in strategic problem framing and formulation. *Organ. Sci., 36*(5), 1962–1979. https://doi.org/10.1287/orsc.2024.19134

Parker, S. K., & Grote, G. (2022). Automation, algorithms, and beyond: Why work design matters more than ever in a digital world. *Appl. Psychol., 71*(4), 1171–1204. https://doi.org/10.1111/apps.12241

Patil, S. V., Srinivas, S. B., Tussing, D. V., & Rhee, J. (2025). Addressing the flexible use of cognitive flexibility constructs: Toward a multifaceted approach. *Acad. Manag. Ann., 19*(1), 74–131. https://doi.org/10.5465/annals.2023.0078

Patwardhan, T., Dias, R., Proehl, E., Kim, G., Wang, M., Watkins, O., … Thacker, P. (2025). *GDPval: Evaluating AI model performance on real-world economically valuable tasks*. arXiv. https://doi.org/10.48550/arXiv.2510.04374.

Pratt, M., & Ashforth, B. (2003). Fostering meaningfulness in working and at work. In K. Cameron, J. Dutton, & R. Quinn (Eds.), *Positive Organizational Scholarship: Foundations of a New Discipline* (pp. 309–327).

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Acad. Manag. Rev., 46*(1), 192–210. https://doi.org/10.5465/amr.2018.0072

Ramaul, L., Ritala, P., Kostis, A., & Aaltonen, P. (2025). Rethinking how we theorize AI in organization and management: A problematizing review of rationality and anthropomorphism. *J. Manag. Stud.*. https://doi.org/10.1111/joms.13246

Ranganathan, A., & Ye, X. M. (2026). AI doesn't reduce work—It intensifies it. *Harv. Bus. Rev.* https://hbr.org/2026/02/ai-doesnt-reduce-work-it-intensifies-it.

Rathje, S., Ye, M., Globig, L. K., Pillai, R. M., De Mello, V. O., & Van Bavel, J. J. (2025). Sycophantic AI increases attitude extremity and overconfidence. *Psy*. https://doi.org/10.31234/osf.io/vmyek_v1. s.

Renner, M., & Chaban, M. A. V. (2025). Real-world gen AI use cases from the world's leading organizations. *Google Cloud Blog*. https://cloud.google.com/transform/101-real-world-generative-ai-use-cases-from-industry-leaders.

Shaw, J. D., Zhu, J., Duffy, M. K., Scott, K. L., Shih, H.-A., & Susanto, E. (2011). A contingency model of conflict and team effectiveness. *J. Appl. Psychol., 96*(2), 391–400. https://doi.org/10.1037/a0021340

Shrestha, Y. R., Ben-Menahem, S. M., & Von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *Calif. Manag. Rev., 61*(4), 66–83. https://doi.org/10.1177/0008125619862257

Stanford University Human-Centered Artificial Intelligence. (2025). *Artificial Intelligence Index Report 2025*.

Steinbach, A. L., Gamache, D. L., & Johnson, R. E. (2019). Don't get it misconstrued: Executive construal-level shifts and flexibility in the upper echelons. *Acad. Manag. Rev., 44*(4), 871–895. https://doi.org/10.5465/amr.2017.0273

Suleyman, M., & Bhaskar, M. (2023). *The coming wave: Technology, power, and the twenty-first century's greatest dilemma*. Crown.

Tey, K. S., Mazar, A., Tomaino, G., Duckworth, A. L., & Ungar, L. H. (2024). People judge others more harshly after talking to bots. *PNAS Nexus, 3*(9), Article pgae397. https://doi.org/10.1093/pnasnexus/pgae397

Van Quaquebeke, N., & Gerpott, F. H. (2023). The now, new, and next of digital leadership: How artificial intelligence (AI) will take over and change leadership as we know it. *J. Leadersh. Org. Stud., 30*(3), 265–275. https://doi.org/10.1177/15480518231181731

Van Quaquebeke, N., & Gerpott, F. H. (2024). Artificial intelligence (AI) and workplace communication: Promises, perils, and recommended policy. *J. Leadersh. Org. Stud., 31*(4), 375–381. https://doi.org/10.1177/15480518241289644

Van Quaquebeke, N., Tonidandel, S., & Banks, G. C. (2025). Beyond efficiency: How artificial intelligence (AI) will reshape scientific inquiry and the publication process. *Leadersh. Q., 101895*. https://doi.org/10.1016/j.leaqua.2025.101895

Wang, Y., Zhu, C., Feng, T., Lu, L. X., & Jia, B. (2026). *Agentic AI and human-in-the-loop interventions: Field experimental evidence from Alibaba's customer service operations (version 1)*. arXiv. https://doi.org/10.48550/ARXIV.2605.14830.

Weick, K. E. (1988). Enacted sense-making in crisis situations. *J. Manag. Stud., 25*(4), 305–317. https://doi.org/10.1111/j.1467-6486.1988.tb00039.x

Xu, F. F., Song, Y., Li, B., Tang, Y., Jain, K., Bao, M., … Neubig, G. (2025). *TheAgentCompany: Benchmarking LLM agents on consequential real world tasks*. arXiv. https://doi.org/10.48550/arXiv.2412.14161.

Zhuo, J., Zhang, S., Fang, X., Duan, H., Lin, D., & Chen, K. (2024). *ProSA: Assessing and understanding the prompt sensitivity of LLMs (version 1)*. arXiv. https://doi.org/10.48550/ARXIV.2410.12405.