

Generative AI in the Workplace: Managerial Implications for Creativity, Productivity, and Job Design

¹Dr. Parul Agrawal, ²Dr Jitendra Patel, ^{3*}Rupam Jyoti Deka

¹Associate Professor, Bharati Vidyapeeth (Deemed to be University) Institute of Management and Research, New Delhi

²Assistant Professor, Stme Nmims Indore

^{3*} Assistant Professor, Bharati Vidyapeeth (Deemed to be University Institute of Management and Research), India

Abstract: Generative artificial intelligence (AI) has moved from laboratory novelty to a general-purpose technology embedded in everyday knowledge work, raising urgent questions for managers about how to capture its benefits while managing its risks. This paper synthesises the emerging empirical and theoretical literature to examine the managerial implications of generative AI across three interrelated domains: productivity, creativity, and job design. Field and laboratory experiments converge on substantial average productivity gains, with effects concentrated among less-experienced workers, suggesting that generative AI can compress performance inequality and act as a skill-levegger. Evidence on creativity is more ambivalent: the technology elevates the novelty and quality of individual outputs yet narrows collective diversity, creating a social dilemma for organisations that depend on variety for innovation. For job design, generative AI reconfigures the task composition of roles along an augmentation–automation spectrum, with consequences for autonomy, skill variety, and meaning that map directly onto established work-design theory. The paper advances an integrative framework and a five-stage managerial approach: assess, redesign, reskill, govern, and learn, arguing that the value of generative AI depends less on the technology itself than on the deliberate organisational choices that surround its adoption.

Keywords: *generative artificial intelligence, productivity, creativity, job design, work design, human–AI collaboration, management*

1. Introduction

The public release of large language models in late 2022 marked an inflexion point in the diffusion of artificial intelligence into the workplace. Within months, tools capable of generating fluent text, functional code, images, and analysis became accessible to hundreds of millions of workers, and organisations across virtually every sector began experimenting with how these systems might be woven into daily work. Unlike earlier waves of automation that targeted routine, codifiable manual and clerical tasks, generative AI reaches into the cognitive and creative core of professional work, drafting, ideation, synthesis, and problem solving (Brynjolfsson, 2022; Eloundou et al., 2023). This breadth is precisely why the technology is best understood as a general-purpose technology whose ultimate economic impact will be shaped by complementary organisational innovation rather than by raw capability alone. For managers, this shift poses a distinctive challenge. The early evidence is genuinely double-edged. Controlled studies report large productivity improvements, yet the same studies caution that benefits are uneven across tasks and workers, that quality can degrade when the technology is applied beyond its competence, and that gains in individual output may come at the cost of homogenised collective work (Dell'Acqua et al., 2023; Doshi & Hauser, 2024). Generative AI is therefore not a turnkey efficiency upgrade but a

sociotechnical intervention whose consequences depend on how work is structured, how people are trained, and how outputs are governed.

This paper synthesises the rapidly accumulating evidence to address a central question: what are the managerial implications of generative AI for creativity, productivity, and job design? These three domains are chosen deliberately. Productivity captures the efficiency case that motivates most adoption decisions. Creativity captures the more contested promise that generative AI augments rather than diminishes human ingenuity. Job design connects both to the lived experience of work, how tasks are bundled into roles, how much autonomy and skill variety people retain, and how meaningful their work remains. Treating these domains in isolation, as much commentary does, obscures the trade-offs that bind them together: the workflow choices that maximise short-run productivity may erode the skill variety and diversity that sustain long-run creativity. The paper proceeds as follows. Section 2 defines generative AI and develops a conceptual framework grounded in work-design theory and the automation–augmentation distinction. Sections 3, 4, and 5 examine the evidence on productivity, creativity, and job design, respectively. Section 6 integrates these strands into managerial implications and a staged implementation framework. Section 7 considers risks, limitations, and directions for future research, and Section 8 concludes.

2. Conceptual Background and Framework

2.1 Defining Generative AI

Generative AI refers to a class of machine-learning systems, most prominently large language models and related multimodal models, that produce novel content such as text, code, images, or audio in response to natural-language prompts. These systems are trained on vast corpora to predict plausible continuations of input, and their defining managerial property is generality: a single model can perform a wide range of tasks without task-specific programming (Eloundou et al., 2023). Estimates of exposure underscore the breadth of potential impact; one influential analysis concluded that a large share of the workforce could have a meaningful fraction of their tasks affected by language-model capabilities, with higher-wage and more educated occupations among the more exposed, an inversion of the pattern associated with prior automation (Eloundou et al., 2023; Felten et al., 2021). Crucially, exposure is not the same as substitution. Whether a capable model displaces a worker, augments one, or leaves a role untouched depends on how organisations choose to deploy it. This indeterminacy is the source of both the opportunity and the managerial responsibility that this paper foregrounds.

2.2 Theoretical Lenses

Three theoretical perspectives structure the analysis. The first is the automation–augmentation distinction. Automation substitutes machine capability for human effort on a task, whereas augmentation combines human and machine capabilities so that each compensates for the other's limitations (Raisch & Krakowski, 2021). Raisch and Krakowski (2021) describe an automation–augmentation paradox: pursuing automation to cut costs can undermine the human expertise on which effective augmentation depends, so the two strategies are in tension over time. This framing is central because generative AI can be deployed at almost any point along the spectrum, and the chosen point has direct consequences for productivity, creativity, and job quality.

The second lens is work-design theory, anchored in the job characteristics model (Hackman & Oldham, 1976). That model holds that motivation, satisfaction, and performance are shaped by five core job characteristics: skill variety, task identity, task significance, autonomy, and feedback. Because generative AI alters the bundle of tasks a worker performs, it directly changes these characteristics. Recent work-design scholarship argues that the rise of algorithms and AI makes intentional work design more important than ever, not less, because technologies that are adopted without redesign tend to strip autonomy and variety from jobs even when that is not the intent (Parker & Grote, 2022).

The third lens concerns the componential nature of creativity. Creativity in organisations is widely modelled as the product of domain-relevant expertise, creativity-relevant processes, and intrinsic motivation, operating within a supportive social environment (Amabile & Pratt, 2016). Generative AI can supply or substitute for components of this system, offering domain knowledge and divergent ideas. However, in doing so, it may also reshape motivation and the diversity of the idea pool, with ambiguous net effects.

Figure 1 integrates these lenses into a single framework. Generative AI capabilities act through three mediating mechanisms: task augmentation and automation, skill levelling and access, and anchoring and homogenization to influence the three outcome domains. Managerial levers shape both how the mechanisms operate and how the outcomes are realised, and organisational feedback closes the loop as firms monitor results and reconfigure work accordingly.




Figure 1. A conceptual framework linking generative AI capabilities, mediating mechanisms, workplace outcomes, and managerial levers. Solid arrows denote primary influences; the dashed arrow denotes the organisational feedback loop.

3. Generative AI and Productivity

3.1 The Empirical Case for Productivity Gains

The strongest evidence on generative AI concerns productivity, and it is unusually consistent across settings. In a preregistered experiment with 453 college-educated professionals completing realistic writing tasks, Noy and Zhang (2023) found that access to ChatGPT reduced average completion time by roughly 40% while raising output quality by about 18%. In one of the first

large-scale field deployments, Brynjolfsson et al. (2025) studied more than 5,000 customer-support agents. They found that access to an AI conversational assistant increased the number of issues resolved per hour by about 15% on average. In software development, an experiment on GitHub Copilot reported that developers using the tool completed a programming task roughly 56% faster than a control group (Peng et al., 2023). Moreover, in a field experiment with 758 management consultants, Dell'Acqua et al. (2023) found that consultants using GPT-4 on tasks within the model's competence completed more tasks, worked faster, and produced output rated more than 40% higher in quality than peers without access. Figure 2 summarises these headline effects. Although the metrics differ, time saved, throughput, and quality are not directly comparable, the direction is uniform, and the magnitudes are large relative to the incremental gains typical of prior information technologies. For managers, the practical implication is that meaningful efficiency improvements are achievable quickly and at low marginal cost for a wide range of cognitive tasks.

What distinguishes these results from earlier automation is their reach into the heart of professional work. Prior generations of office technology accelerated specific, well-bounded operations such as calculation, storage, and communication. Generative AI, by contrast, contributes to open-ended cognitive activities, composing arguments, generating code, and synthesising disparate information that were previously considered resistant to mechanisation (Brynjolfsson, 2022; Eloundou et al., 2023). The mechanism behind the measured gains is also instructive. In writing tasks, the technology substituted for effort on rough drafting and shifted human attention toward higher-order activities (Noy & Zhang, 2023), while in customer support, it functioned by surfacing the effective practices of experienced agents to their less-experienced peers (Brynjolfsson et al., 2025). In each case, the productivity gain reflects a reallocation of human effort rather than the wholesale removal of the human from the task, an important signal that the dominant near-term mode is augmentation rather than full automation.




Figure 2. Measured productivity effects of generative AI across representative field and laboratory studies. Metrics differ across studies; bars denote the direction and magnitude of the primary reported effect rather than a common unit.

3.2 The Jagged Frontier and the Limits of Gains

These averages conceal an important boundary condition. Dell'Acqua et al. (2023) introduced the metaphor of a "jagged technological frontier" to capture the uneven contour of model capability: tasks that appear similar in difficulty can fall on opposite sides of what the model does well. On tasks inside the frontier, AI assistance improved performance markedly; on tasks outside it, the same assistance led consultants to perform substantially worse than those without AI, because workers tended to accept fluent but flawed output uncritically. This finding reframes the managerial challenge from "how much will AI help?" to "where does AI help, and how do we keep people engaged enough to catch its errors?" The jagged frontier also explains why productivity gains do not translate automatically into organisational value. Faster production of low-value or erroneous work is not progress, and overreliance on confident but incorrect output can introduce risk into decisions that were previously slower but sounder. Capturing genuine productivity, therefore, requires investment in verification, prompt design, and the cultivation of judgment about when to trust the tool, complementary practices that the technology does not provide on its own.

3.3 Heterogeneous Effects and Skill Levelling

A robust and managerially consequential pattern is that productivity gains are distributed unevenly across workers. Brynjolfsson et al. (2025) found that the largest improvements on the order of 34% accrued to the least experienced and lowest-skilled agents, while the most experienced agents saw small gains and, in some respects, slight declines in quality. Noy and Zhang (2023) similarly observed that generative AI compressed the productivity distribution by helping lower-ability

workers most, thereby reducing inequality between workers. Figure 3 depicts this skill-levelling pattern.




Figure 3. The skill-levelling pattern. Stylised representation of the finding that productivity gains from generative AI are largest for less-experienced workers and smallest for the most experienced, based on Brynjolfsson et al. (2025) and Noy and Zhang (2023).

The mechanism appears to be diffusion of tacit expertise: the model encodes and distributes the practices of high performers, allowing novices to move down the experience curve more quickly (Brynjolfsson et al., 2025). For managers, skill levelling carries mixed implications. It can shorten onboarding, raise the floor of performance, and expand who can perform skilled work. However, it may also compress wage premia tied to experience, complicate the rationale for traditional career ladders, and raise the question of how organisations will continue to develop the deep expertise that AI itself depends upon as a training signal. If novices rely on AI rather than building expertise, the pipeline of experts who can recognise the model's errors may thin over time, another expression of the automation–augmentation paradox (Raisch & Krakowski, 2021).

3.4 From Individual Gains to Organisational Value

A final caution concerns the leap from individual task-level gains to organisational performance. The experimental evidence measures the productivity of individuals on circumscribed tasks, but firm value depends on coordination, process design, and the conversion of saved time into useful output. Time saved by one worker can be dissipated if downstream processes cannot absorb the additional throughput, if freed hours are reallocated to low-value activity, or if accelerated output shifts a bottleneck elsewhere in the workflow. History suggests that the productivity payoff of general-purpose technologies materialises only after firms reorganise their processes around them, often with a substantial lag (Brynjolfsson, 2022). The managerial task, therefore, is not merely to grant access to capable tools but to redesign the surrounding workflows so that individual gains aggregate into organisational performance. This theme connects productivity directly to job design.

4. Generative AI and Creativity

4.1 Augmenting Individual Creativity

Creativity has long been considered a distinctively human capacity, which makes generative AI's incursion into creative work both striking and contested. The evidence suggests that, at the level of the individual, generative AI can enhance creative output. In a controlled experiment in which writers could draw on AI-generated story ideas, Doshi and Hauser (2024) found that access to such ideas caused stories to be rated as more novel, better written, and more enjoyable, with the largest gains among writers who were less creative to begin with. This mirrors the skill-leveiling pattern seen in productivity: AI raises the floor, narrowing the gap between more and less talented individuals. From a componential view of creativity, the model supplies domain-relevant material and divergent prompts that less-expert individuals would struggle to generate unaided (Amabile & Pratt, 2016). For managers, this implies that generative AI can broaden participation in creative tasks, such as idea generation, design exploration, marketing copy, and early-stage R&D, beyond a small set of acknowledged specialists. Brainstorming and option generation, in particular, are domains where the model's capacity to produce many candidate ideas quickly can expand the search space a team considers.

It is useful, however, to distinguish among the phases of creative work, because generative AI does not contribute equally to all of them. In the divergent phase, generating a wide range of candidate ideas, the technology is a powerful complement, supplying volume and variety that can break individuals out of habitual patterns. In the convergent phase, evaluating, selecting, and combining ideas into a coherent and contextually appropriate solution, the human contribution remains decisive, both because judgment about fit and value depends on tacit knowledge of the audience and purpose and because the model has no genuine stake in the outcome. The most promising creative configurations, therefore, assign the model to expand the option space while reserving discrimination and synthesis for people. This allocation also helps guard against the homogenization risk discussed next.

4.2 The Homogenization Risk

The same study delivers a cautionary counterpoint. While individual stories became more creative, the body of AI-assisted stories became more similar to one another than those produced by unaided humans (Doshi & Hauser, 2024). Generative AI thus raises individual creativity while reducing collective diversity. Doshi and Hauser (2024) characterise this as a social dilemma: each writer is individually better off drawing on the model, but the aggregate result is a narrower range of novel content. Because models are trained on common data and tend toward statistically typical outputs, heavy reliance can anchor users on similar ideas and erode the variety on which genuine innovation depends. This dynamic is consequential for organisations precisely because innovation is a collective, portfolio-level phenomenon. A firm whose teams all converge on the same AI-suggested concepts may produce competent but undifferentiated products, weakening competitive distinctiveness. The risk is amplified at the market level: if competitors use similar models in similar ways, entire industries may drift toward sameness. Managers concerned with long-run innovation, therefore, cannot evaluate creative AI tools solely by their effect on individual output; they must also attend to diversity at the team and organisational level.

4.3 Designing for Creative Augmentation

Reconciling these findings points toward a deliberate model of human–AI co-creation rather than substitution. Practices that appear to preserve diversity and motivation include using AI for divergent idea generation while reserving selection, combination, and refinement for human judgment; deliberately prompting for unconventional or contrarian options; and treating model output as raw material to be transformed rather than a finished product to be accepted. Maintaining intrinsic motivation matters as well: if workers experience creative tasks as merely editing machine output, the engagement that drives original work may decline (Amabile & Pratt, 2016). The managerial goal is to position generative AI as a stimulus that expands human creative range, not as an anchor that quietly narrows it. At the portfolio level, organisations can take more structural steps to protect variety. These include diversifying the tools, prompts, and source materials that teams draw on so that creative work does not all flow through a single model's statistical centre of gravity; preserving channels for unaided human ideation alongside AI-assisted work; and incorporating explicit measures of distinctiveness into how creative output is evaluated, so that originality is rewarded rather than penalised by efficiency metrics. Framing the issue this way reframes creativity from a purely individual attribute into an organisational capability that managers can actively cultivate or inadvertently erode through their design of the creative process.

5. Generative AI and Job Design

5.1 Reconfiguring the Task Composition of Roles

Generative AI rarely automates an entire occupation; instead, it automates or augments particular tasks within a job (Acemoglu & Restrepo, 2019; Eloundou et al., 2023). Because jobs are bundles of tasks, shifting which tasks are performed by people changes the character of the role. Noy and Zhang (2023) observed this directly: ChatGPT did not simply speed up writing but restructured it, shifting effort away from rough drafting and toward idea generation and editing. Multiplied across many roles, such task reallocation is the central job-design consequence of generative AI. Figure 4 situates this reconfiguration along the augmentation–automation spectrum. At one end lie routine, codifiable tasks, formatting, summarising, and boilerplate drafting that AI can largely perform, freeing human time. In the middle lie hybrid tasks of co-production, where AI generates and the human directs, edits, and judges. At the other end lie judgment-intensive tasks, strategy, ethics, relationship building, and decisions requiring tacit and contextual knowledge, where human capability remains primary. The job-design question is how to recompose roles so that capacity freed at the automatable end is reinvested in more valuable and motivating work rather than absorbed as a higher pace of low-value output.




Figure 4. The augmentation–automation spectrum and its job-design implications. Tasks are distributed across a continuum from full automation to pure human judgment, with hybrid co-production in between; the design challenge is to reinvest freed capacity in higher-value work.

Reconfiguration also operates above the level of the individual role. When the boundaries of what a single person can accomplish expand, the rationale for dividing labour among specialists may weaken, allowing roles to be broadened or recombined. A marketer who can draft passable code, or an analyst who can produce polished prose, occupies a different position in the division of labour than before. This can enrich jobs by increasing their scope and task identity. However, it can also blur professional boundaries and raise coordination questions about who is accountable for work that now spans former specialities. Managers should therefore treat job design as a question not only of individual task allocation but of how reconfigured roles fit together into a coherent and accountable system of work.

5.2 Implications for Job Characteristics

These shifts map directly onto the job characteristics model (Hackman & Oldham, 1976). Their net effect on motivation depends on how the redesign is handled. On the positive side, offloading tedious tasks can increase the share of time spent on stimulating, higher-order work, potentially raising skill variety and task significance; consistent with this, Noy and Zhang (2023) reported increases in job satisfaction and self-efficacy among workers using the tool. On the negative side, several risks threaten job quality. Autonomy can erode if AI systems prescribe how work is done or if monitoring intensifies. Skill variety can narrow if interesting tasks are automated and workers are left to verify machine output, a shift toward a supervisory, correction-oriented role that some find less engaging. Task identity can blur when a worker contributes only fragments to an AI-assembled whole. Moreover, feedback can become distorted if the model, rather than the work itself or other people, becomes the primary source of cues. Parker and Grote (2022) argue that these are not inevitable outcomes of the technology but consequences of design choices. Organisations that adopt AI without rethinking roles tend to default to efficiency-maximising configurations that quietly degrade autonomy and variety. In contrast, those who treat work design as a deliberate variable can use the same technology to enrich jobs. The implication is that managers should treat job redesign as a primary lever rather than an afterthought.

5.3 Skill Demands and Role Redefinition

As task composition shifts, so do the skills that roles demand. The relative value of routine drafting and information retrieval declines, while the value of skills that complement the technology rises: framing problems and prompts effectively, critically evaluating and verifying output, exercising judgment about when to rely on the model, and integrating AI contributions with human and organisational context (Dell'Acqua et al., 2023). New hybrid roles are emerging that pair domain expertise with fluency in directing AI systems. For managers, this entails redefining role descriptions, performance criteria, and development pathways so that they reward the distinctively human and AI-complementary capabilities that now carry the most value.

6. Integrative Managerial Implications

The preceding sections reveal a consistent meta-lesson: the value generative AI creates depends less on the technology than on the organisational choices surrounding it. The same model can raise

or lower quality, enrich or degrade jobs, and expand or narrow creative diversity, depending on how tasks are allocated, how people are trained, and how outputs are governed. This section consolidates the implications into a staged managerial approach, summarised in Figure 5.




Figure 5. A five-stage managerial framework for integrating generative AI: assess, redesign, reskill, govern, and learn, connected by a continuous feedback loop.

6.1 Assess: Map the Jagged Frontier

Effective adoption begins with a clear-eyed assessment of where generative AI genuinely helps and where it harms. Because capability is jagged, blanket mandates to "use AI everywhere" are counterproductive (Dell'Acqua et al., 2023). Managers should identify, at the task level, which activities fall reliably inside the model's competence, which require careful human verification, and which should remain human-led. This mapping is best developed empirically through structured piloting rather than assumed from vendor claims.

6.2 Redesign: Rebuild Jobs, Not Just Tasks

Once the task mix is understood, roles should be intentionally recomposed. The objective is to reinvest freed capacity in higher-value, more motivating work and to protect the core job characteristics that sustain engagement (Hackman & Oldham, 1976; Parker & Grote, 2022). In practice, this means designing jobs in which workers retain meaningful autonomy and decision rights, preserving variety by ensuring that automation does not strip out all the engaging tasks, and keeping humans in roles of genuine judgment rather than mere output verification.

6.3 Reskill: Build AI-Complementary Capabilities

Realising value requires investment in the human capabilities that complement the technology. These include AI literacy and prompt design, but more importantly, the critical-evaluation skills needed to catch the model's confident errors and the domain judgment needed to know when to trust it (Dell'Acqua et al., 2023). Because skill levelling can weaken traditional expertise-development pathways, organisations must take deliberate steps to continue cultivating deep expertise, ensuring that reliance on AI among novices does not hollow out the pipeline of experts on whom effective oversight depends (Raisch & Krakowski, 2021).

6.4 Govern: Establish Guardrails and Trust

Generative AI introduces risks of factual errors, bias, confidentiality exposure, and intellectual-property ambiguity that require governance. Managers should set clear standards for quality

assurance and human review of consequential outputs, define where AI use is and is not appropriate, and be transparent with workers and customers about how the technology is used. Trust is a two-sided concern: employees must trust that AI is being introduced to augment rather than replace them, or adoption will meet resistance, while organisations must avoid the overtrust that the jagged frontier punishes.

6.5 Learn: Monitor, Iterate, and Protect Diversity

Finally, because the technology and its effects are evolving quickly, integration should be treated as a continuous learning process rather than a one-time deployment. Organisations should measure not only efficiency but also quality, job satisfaction, and, given the homogenization risk, the diversity of creative output, iterating on workflows as they learn (Doshi & Hauser, 2024). The feedback loop in Figure 5 emphasises that assessment, redesign, reskilling, and governance are not sequential milestones to be completed once but practices to be revisited as capabilities and conditions change.

7. Challenges, Limitations, and Future Research

Several challenges qualify the optimistic productivity narrative and define an agenda for research. First, much of the strongest evidence comes from controlled experiments and short-run field studies; the durability of productivity gains, their effect on firm-level outcomes, and the long-run consequences of skill levelling for career development and wage structures remain open questions. Aggregate productivity statistics have so far shown only modest movement, echoing the historical lag between the arrival of general-purpose technologies and their measurable economic payoff as complementary organisational innovations diffuse (Brynjolfsson, 2022). Second, the homogenization finding raises a research question of real economic weight: under what conditions does widespread AI use erode collective diversity, and which organisational practices counteract it (Doshi & Hauser, 2024)? Third, the work-design consequences of generative AI for autonomy, meaning, and well-being are understudied relative to its productivity effects, even though they may prove more decisive for long-run performance and retention (Parker & Grote, 2022). Fourth, distributional and ethical questions of who captures the gains, how bias propagates, and how accountability is assigned for AI-assisted decisions demand attention from both researchers and managers. Finally, this paper is a synthesis of a fast-moving and still-thin evidence base; its framework should be read as a structured set of propositions to be tested rather than settled conclusions. As models advance toward more autonomous, agentic capabilities, the boundary between augmentation and automation will shift, and the managerial questions examined here will require continual reassessment.

8. Conclusion

Generative AI presents managers with a genuine opportunity and an equally genuine set of trade-offs. The evidence indicates substantial productivity gains, especially for less-experienced workers; a creativity effect that lifts individuals while threatening collective diversity; and a reconfiguration of tasks that can either enrich or degrade jobs depending on how roles are redesigned. Across all three domains, the recurring lesson is that outcomes are not determined by the technology but by managerial choices about how it is integrated into work. The framework advanced here assesses, redesigns, reskills, governs, and learns, treating generative AI as a

sociotechnical intervention requiring deliberate organisational design rather than a self-executing efficiency tool. Managers who attend only to short-run productivity risk eroding the diversity, expertise, and engagement on which long-run value depends, while those who manage the technology thoughtfully can capture its benefits while sustaining the human capabilities that remain its essential complement. The defining managerial task of this era is not to deploy generative AI but to design the work, the skills, and the safeguards that allow people and machines to be more than the sum of their parts.

References

1. Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labour. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
2. Amabile, T. M., & Pratt, M. G. (2016). The dynamic componential model of creativity and innovation in organisations: Making progress, making meaning. *Research in Organizational Behavior*, 36, 157–183. <https://doi.org/10.1016/j.riob.2016.10.001>
3. Brynjolfsson, E. (2022). The Turing trap: The promise and peril of human-like artificial intelligence. *Daedalus*, 151(2), 272–287. https://doi.org/10.1162/daed_a_01915
4. Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
5. Dell’Acqua, F., McFowland, E., III, Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Krayner, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality (Working Paper No. 24-013). Harvard Business School. <https://doi.org/10.2139/ssrn.4573321>
6. Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), eadn5290. <https://doi.org/10.1126/sciadv.adn5290>
7. Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models (arXiv:2303.10130). arXiv. <https://doi.org/10.48550/arXiv.2303.10130>
8. Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217. <https://doi.org/10.1002/smj.3286>
9. Hackman, J. R., & Oldham, G. R. (1976). Motivation through the design of work: Test of a theory. *Organizational Behavior and Human Performance*, 16(2), 250–279. [https://doi.org/10.1016/0030-5073\(76\)90016-7](https://doi.org/10.1016/0030-5073(76)90016-7)
10. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
11. Parker, S. K., & Grote, G. (2022). Automation, algorithms, and beyond: Why work design matters more than ever in a digital world. *Applied Psychology: An International Review*, 71(4), 1171–1204. <https://doi.org/10.1111/apps.12241>

12. Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot (arXiv:2302.06590). arXiv. <https://doi.org/10.48550/arXiv.2302.06590>
13. Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://doi.org/10.5465/amr.2018.0072>