

The Role of Generative Artificial Intelligence in Modern Business Decision-Making: Applications, Opportunities, and Future Challenges

Rajidi Rammohan Reddy¹, Vinodray Thumar², Amar Jyoti Borah³, T. Vijayakumar⁴, Shabina⁵, Tara Sasanka, C⁶

¹Department Of Management Studies, Trinity College of Engineering and Technology, Peddapalli, Telangana, dr.rajidi@gmail.com, Orcid: 0009-0001-8383-7679

²Computer Engineering Department, Vishwakarma Government Engineering College, GTU, Ahmedabad, Gujarat, India

³Department Of Commerce, The Assam Royal Global University. Kamrup (M), Guwahati, Assam, Amar.borah2015@gmail.com

⁴Department Of Artificial intelligence and machine learning, Dr.Mahalingam College of Engineering and Technology, Coimbatore, Pollachi, Tamil Nadu ORC id:0009-0003-0896-1412

⁵Department of Computer Science, Govt. College Mohali, SAS Nagar, Punjab, dr.shabina89@gmail.com

⁶Mechanical Engineering, R.V.R & J.C. College of Engineering, Guntur, Andhra Pradesh, tarasasankac@gmail.com

Abstract: Generative artificial intelligence (GenAI) has moved rapidly from an experimental technology to a tool actively reshaping how managers gather information, evaluate options, and reach decisions across strategic, operational, and customer-facing business functions. This paper reviews the theoretical and empirical literature on GenAI's role in business decision-making, situating recent large language model (LLM)-based applications within the older organizational decision-making and bounded-rationality literature that predates generative AI by decades. The review synthesizes controlled productivity experiments, large-scale economic-potential estimates, organizational decision-structure theory, and the emerging literature on human-AI complementarity and its limits, including evidence that AI assistance can degrade performance when applied outside a model's effective capability frontier. Particular attention is given to the distinction between GenAI's demonstrated value in well-structured, information-synthesis-heavy decision tasks and its more contested role in tasks requiring novel judgment or accountability. Comparative tables summarize reported productivity effects, business function applications, and organizational barriers to adoption across the reviewed literature. The paper concludes that GenAI's current business value is concentrated in augmenting, rather than automating, decision-making, with the strongest evidence for productivity gains among relatively lower-skilled or lower-performing decision-makers, and identifies the mapping of AI capability boundaries within specific decision domains as the central future research prospect. Generative artificial intelligence (GenAI) has moved rapidly from an experimental technology to a tool actively reshaping how managers gather information, evaluate options, and reach decisions across strategic, operational, and customer-facing business functions. This paper reviews the theoretical and empirical literature on GenAI's role in business decision-making, situating recent large language model (LLM)-based applications within the older organizational decision-making and bounded-rationality literature that predates generative AI by decades. The review synthesizes controlled productivity experiments, large-scale economic-potential estimates, organizational decision-structure theory, and the emerging literature on human-AI complementarity and its limits, including evidence that AI assistance can degrade performance when applied outside a model's effective capability frontier. Particular attention is given to the distinction between GenAI's demonstrated value in well-structured, information-synthesis-heavy decision tasks and its more contested role in tasks requiring novel judgment or accountability. Comparative tables summarize reported productivity effects, business function applications, and organizational barriers to adoption across the reviewed literature. The paper concludes that GenAI's current business value is concentrated in augmenting, rather than automating, decision-making, with the strongest evidence for productivity gains among relatively lower-skilled or lower-performing decision-makers, and identifies the mapping of AI capability boundaries within specific decision domains as the central future research prospect.

Keywords: Generative Artificial Intelligence, Business Decision-Making, Large Language Models, Organizational Behavior, Human-AI Collaboration, Productivity, Bounded Rationality.



1. Introduction

Organizational decision-making has long been understood as a process constrained by cognitive and informational limits rather than one approximating the fully rational, comprehensive-search model assumed by classical economic theory. Simon's foundational theory of bounded rationality established that managers do not optimize across the full space of available alternatives but instead "satisfice," selecting the first option that meets an acceptable threshold, a behavior Simon attributed directly to the practical impossibility of gathering and processing all relevant information within the time and cognitive resources any real decision-maker has available [1]. Kahneman's subsequent dual-process account of human judgment extended this bounded-rationality framework, distinguishing a fast, intuitive, heuristic-driven mode of reasoning from a slower, more deliberate analytical mode, and demonstrating across decades of experimental evidence that even expert decision-makers routinely default to the faster, heuristic mode in ways that introduce systematic, predictable biases into business and financial judgment [2]. This body of theory, developed well before the emergence of generative AI, frames the central question this paper addresses: does GenAI's capacity to rapidly synthesize, summarize, and generate information at scale address the specific cognitive and informational constraints that bounded-rationality theory identifies as the root cause of imperfect business decision-making, or does it introduce a distinct set of limitations of its own [1], [2]

Agrawal, Gans, and Goldfarb's influential reframing of artificial intelligence's economic function characterized AI, including the generative models that emerged several years after their initial formulation, primarily as a prediction technology, a tool that reduces the cost of generating predictions about uncertain future states, and argued that cheaper prediction should be expected to increase the value of complementary human judgment, the capacity to weigh a prediction's payoff structure and decide on an appropriate action, rather than eliminating the need for such judgment [12]. This economic framing anticipated the subsequent empirical literature's central finding: generative AI tools have consistently demonstrated their strongest, most reproducible value in prediction- and synthesis-heavy decision-support tasks, while the judgment component of business decision-making has remained comparatively resistant to full automation [3], [4], [13].

Brynjolfsson, Li, and Raymond's large-scale field experiment, conducted within a customer support function at a Fortune 500 software company, provided direct empirical evidence for this pattern, finding that access to a generative AI conversational assistant increased customer support agents' resolved-issues-per-hour productivity by roughly 14 percent on average, with the effect concentrated overwhelmingly among newer and lower-skilled agents, whose productivity gains substantially exceeded those of experienced, high-performing agents, whose output the AI tool barely affected [3]. This skill-distribution finding closely parallels Noy and Zhang's earlier controlled experiment on professional writing tasks, which similarly found that generative AI's largest productivity gains accrued to relatively lower-performing participants, a pattern the authors interpret as evidence that generative AI functions partly as a knowledge-leveling technology that compresses, rather than uniformly amplifies, the existing performance distribution among decision-makers and knowledge workers [11].

Dell'Acqua et al.'s large-scale field experiment conducted with Boston Consulting Group consultants introduced an important qualification to this generally positive productivity picture, finding that consultants using GPT-4 for tasks within the model's effective capability frontier significantly outperformed a control group on both speed and quality, while consultants who used the same tool for a task specifically designed to fall outside the model's effective capability, a task requiring a type of business judgment the model could not reliably perform, underperformed the control group, having been led astray by the AI system's confidently presented but substantively incorrect output [4]. This finding, widely referred to subsequently as the "jagged technological frontier," established that generative AI's business value is not uniform across decision task types but depends critically on whether a given task falls within or outside the model's actual, rather than perceived, capability boundary, a distinction with direct implications for how organizations should structure AI-assisted decision workflows [4].

This paper synthesizes the organizational decision-theory, controlled-experiment, and economic-potential literatures into a single comparative framework, examining GenAI's application across specific business functions, the organizational and governance barriers constraining broader adoption, and the emerging evidence on where human-AI complementarity in decision-making succeeds and where it breaks down.

Research Objectives

- To review the bounded-rationality and economic-prediction theoretical foundations relevant to evaluating generative AI's role in business decision-making.
- To synthesize controlled experimental and field evidence on generative AI's productivity effects across business functions.
- To examine organizational decision-structure theory and its application to AI-augmented decision-making.
- To evaluate the boundary conditions under which generative AI assistance improves versus degrades business decision quality.
- To identify future research priorities for mapping AI capability frontiers within specific business decision domains.

2. Literature Review

2.1 Bounded Rationality and the Economics of Prediction

Simon's satisficing model established that organizational decision-making is fundamentally a search process operating under severe informational and cognitive constraints, and that improving the speed or breadth of information available to a decision-maker should, in principle, expand the set of alternatives a decision-maker can meaningfully evaluate before selecting a satisfactory option [1]. Agrawal, Gans, and Goldfarb's prediction-machine framework operationalized this insight for the AI era specifically, arguing that as the cost of prediction falls, decision-making tasks decompose usefully into a prediction component, well-suited to AI automation, and a judgment component, the assignment of value or payoff to different predicted outcomes, that remains a distinctly human function requiring accountability, contextual understanding, and, frequently, organizational authority that current AI systems do not possess [12]. Kahneman's dual-process theory adds a further dimension to this framework: since human judgment itself is subject to systematic heuristic biases, a well-designed AI-prediction system has the theoretical potential not merely to supply information faster but to counteract specific, well-documented human decision biases, provided the AI system's own outputs do not introduce comparable or novel biases of their own [2].

2.2 Organizational Decision Structures and AI Integration

Shrestha, Ben-Menahem, and von Krogh's analysis of organizational decision-making structures in the age of artificial intelligence proposed a typology distinguishing full human control, hybrid human-AI decision-making, and full AI automation of a given decision, arguing that the appropriate structure for a given organizational decision depends on the decision's frequency, its stakes, and the degree to which the underlying prediction task is well specified, with hybrid structures recommended for the large majority of consequential business decisions given current AI capability [8]. Von Krogh's broader theoretical treatment of artificial intelligence in organizations argued that AI's integration into organizational decision-making constitutes a genuinely novel phenomenon warranting new theory rather than a straightforward extension of prior information-technology adoption research, since generative and predictive AI systems, unlike earlier decision-support software, actively generate novel content and recommendations rather than merely retrieving or structuring pre-existing organizational data [9]. Berente, Gu, Recker, and Santhanam's subsequent framework for managing artificial intelligence in organizations identified the opacity of many AI systems' internal reasoning, including generative language models, as a distinct organizational management challenge, since managers integrating AI into decision workflows must establish appropriate levels of trust and oversight without the ability to fully audit or explain the system's underlying reasoning process in the way they could audit a human subordinate's stated rationale [10].

2.3 Field Experimental Evidence on Productivity

Brynjolfsson, Li, and Raymond's customer-support field experiment provides the most methodologically rigorous large-scale field evidence currently available on generative AI's business productivity effect, benefiting from access to actual firm performance data across thousands of agents rather than a laboratory task, and found that the AI tool's productivity benefit operated primarily by helping less experienced agents apply organizational best practices more consistently, effectively accelerating the on-the-job learning curve that would otherwise require months of accumulated experience [3]. Dell'Acqua et al.'s BCG consultant experiment, while conducted with a professional population further up the skill and compensation distribution than Brynjolfsson, Li, and Raymond's customer-support agents, similarly found substantial productivity and quality gains for tasks within the model's capability frontier, with consultants completing 12.2 percent more tasks on average and producing work rated 40 percent higher in quality by

blind evaluators when using GPT-4 for appropriately matched tasks [4]. Peng et al.'s controlled study of software developers using GitHub Copilot, an AI coding assistant built on generative language model technology, found that developers using the tool completed a standardized programming task 55.8 percent faster than a control group not using the tool, extending the productivity evidence base from customer service and consulting specifically into software engineering, a technical decision-making and problem-solving domain with its own distinct evaluation criteria [6].

2.4 Macroeconomic and Cross-Functional Potential Estimates

Chui et al.'s McKinsey Global Institute analysis of generative AI's economic potential estimated that generative AI technologies could add trillions of dollars in annual value to the global economy, with the largest concentrations of estimated value in customer operations, marketing and sales, software engineering, and research and development functions, a cross-functional distribution the authors attribute to these functions' relatively high proportion of language-based, information-synthesis-heavy tasks amenable to current generative AI capability [5]. Eloundou, Manning, Mishkin, and Rock's labor-market exposure analysis complemented this function-level estimate with an occupation-level analysis, finding that approximately 80 percent of the U.S. workforce could have at least 10 percent of their work tasks affected by large language models, with higher-wage, more information-processing-intensive occupations generally showing greater exposure than lower-wage, physically-intensive occupations, a pattern with direct implications for which layers of business decision-making, typically concentrated among higher-wage information-processing roles, are most likely to be affected by continued generative AI adoption [7]. Bughin et al.'s earlier McKinsey Global Institute analysis of AI more broadly, conducted before the generative AI wave specifically, had already identified marketing and sales, supply chain management, and product development as functions with disproportionately high AI value potential, a finding that the subsequent generative-AI-specific estimates from Chui et al. largely reaffirmed and extended rather than contradicted [5], [14].

2.5 The Boundary of Human-AI Complementarity

Dell'Acqua et al.'s capability-frontier finding, that AI assistance improved performance within the model's effective capability but degraded performance when consultants applied the same tool to a task specifically designed to exceed that capability, represents the field's most direct empirical evidence that generative AI's business decision-making value is not unconditionally positive [4]. The authors documented that underperforming consultants who used the AI tool outside its capability frontier were more likely to accept the model's confidently stated but incorrect output without adequate independent verification, a pattern the authors term a false sense of confidence effect, closely related to the automation-bias phenomenon documented in earlier human-automation interaction research predating generative AI specifically [4]. Davenport and Ronanki's earlier survey-based study of AI implementation across nearly 250 organizations, conducted before the generative AI wave but examining structurally similar cognitive-automation technologies, found that organizations reporting the most successful AI implementations were disproportionately those that had explicitly redesigned decision workflows around AI's specific strengths and limitations, rather than simply inserting an AI tool into an existing, unmodified decision process, a finding directly relevant to the capability-frontier boundary Dell'Acqua et al. later documented empirically for generative AI specifically [4], [13].

2.6 Ethical, Governance, and Adoption Barriers

Floridi et al.'s AI4People framework articulated a set of ethical principles, beneficence, non-maleficence, autonomy, justice, and explicability, intended to guide responsible AI deployment across sectors including business, and specifically emphasized explicability, the capacity to understand and articulate why an AI system produced a given output, as a principle with direct relevance to business decision-making contexts where decision-makers may bear legal or fiduciary accountability for AI-informed decisions they cannot fully explain [15]. Susarla, Gopal, Thatcher, and Sarker's Janus-faced characterization of generative AI's organizational impact, examined in the broader information-systems literature, argued that the same generative capability accelerating routine business content production and decision-support tasks simultaneously introduces governance risks around misinformation, intellectual property ambiguity, and overreliance on unverified AI-generated business analysis entering consequential organizational decisions [16]. Sætra's analysis of generative AI's longer-term organizational and societal trajectory argued that the technology's ultimate business value will depend substantially on governance and implementation choices made in the near term, rather than on the technology's raw capability alone, a conclusion consistent with Davenport and Ronanki's earlier empirical finding that successful AI implementation depends heavily on deliberate workflow redesign [13], [17].

2.7 Research Gap

While the productivity literature reviewed in Sections 2.3 and 2.4 has established robust, methodologically diverse evidence of generative AI's aggregate business value, and Dell'Acqua et al.'s capability-frontier finding has established that this value is conditional rather than unconditional, comparatively few studies directly map the specific boundary between well-suited and poorly-suited decision tasks across the range of business functions Chui et al. and Eloundou et al. identify as high-exposure, leaving organizations with limited task-level guidance for the workflow-redesign approach Davenport and Ronanki identify as central to successful implementation [3], [4], [5], [7], [13]. This review addresses this gap by organizing the reviewed evidence within a comparative framework spanning business function, reported productivity effect, and capability-frontier risk.

3. Methodology

This study adopts a qualitative, descriptive, and comparative research design based on synthesis of peer-reviewed field experiments, organizational-theory literature, and industry economic-potential analyses addressing generative AI's role in business decision-making. Reviewed studies were compared along four axes: business function examined (customer operations, consulting/strategy, software engineering, or cross-functional); evidence type (controlled field experiment, laboratory experiment, survey-based, or macroeconomic estimation); reported productivity or quality effect and its distribution across participant skill level; and whether the study addresses capability-frontier or task-suitability boundaries explicitly.

4. Results and Discussion

4.1 Reported Productivity Effects Across Business Functions

Table 1 summarizes the reported productivity and quality effects of the major field and laboratory studies reviewed, organized by business function rather than publication chronology.

Table 1. Reported Generative AI Productivity Effects by Business Function

Business Function	Study	Study Design	Reported Effect	Skill Distribution of Benefit
Customer support/operations	Brynjolfsson, Li & Raymond [3]	Large-scale field experiment (Fortune 500 firm)	~14% increase in issues resolved per hour	Concentrated among newer/lower-skilled agents
Management consulting	Dell'Acqua et al. [4]	Field experiment (BCG consultants)	+12.2% tasks completed; +40% quality rating (in-frontier tasks)	Broad gains in-frontier; losses out-of-frontier
Software engineering	Peng et al. [6]	Controlled experiment (GitHub Copilot)	55.8% faster task completion	Not disaggregated by skill in original study
Professional writing	Noy & Zhang [11]	Controlled experiment	Reduced task time, improved quality	Concentrated among lower-baseline performers

4.2 Cross-Functional Economic Potential and Exposure

Table 2 presents the macroeconomic-scale evidence reviewed in Section 2.4, distinct in scope and method from the individual-level productivity studies in Table 1, contrasting estimated value potential against workforce exposure by function.

Table 2. Estimated Economic Potential and Workforce Exposure by Business Function

Business Function	Estimated Value Potential (Chui et al. [5])	Underlying Driver	Workforce Exposure Note (Eloundou et al. [7])

Customer operations	High	High volume of language-based synthesis and response tasks	High-exposure occupational category
Marketing and sales	High	Content generation, personalization at scale	High-exposure occupational category
Software engineering	High	Code generation, debugging assistance	High-exposure occupational category
Research and development	High	Literature synthesis, hypothesis generation support	Moderate-to-high exposure
Physically-intensive operations (e.g., manufacturing floor tasks)	Lower (relative)	Limited applicability of language-based generative tools	Low-exposure occupational category

4.3 A Task-Suitability Risk Matrix

Unlike Tables 1 and 2, which summarize reported effects, Table 3 organizes the capability-frontier and organizational-governance evidence reviewed in Sections 2.5 and 2.6 into a risk-oriented matrix, classifying business decision-task types by their apparent suitability for current generative AI assistance based on the reviewed evidence.

Table 3. Task-Suitability Risk Matrix for Generative AI-Assisted Business Decisions

Task Characteristic	Example Business Decision Task	Suitability for Current GenAI Assistance	Supporting Evidence
Well-defined synthesis/retrieval task within known domain	Summarizing customer complaint patterns; drafting standard reports	High suitability	Brynjolfsson, Li & Raymond [3]; Chui et al. [5]
Routine code generation/debugging within established codebase	Writing boilerplate code, unit tests	High suitability	Peng et al. [6]
Novel strategic judgment requiring accountability	Business-model pivot recommendation outside model's trained domain	Low suitability; risk of false confidence	Dell'Acqua et al. [4]
High-stakes decision with legal/fiduciary accountability	Regulatory compliance determination, resource-allocation sign-off	Requires explicability; governance risk	Floridi et al. [15]; Berente et al. [10]
Decision embedded in un-redesigned legacy workflow	Inserting AI output directly into existing approval process without adaptation	Elevated implementation-failure risk	Davenport & Ronanki [13]

4.4 Discussion

Taken together, the evidence in Tables 1 through 3 supports a conditional rather than unconditional assessment of generative AI's value in business decision-making, directly consistent with the bounded-rationality and prediction-machine theoretical frameworks reviewed in Section 2.1. Table 1's productivity findings, while methodologically diverse in design and population, converge on a consistent skill-distribution pattern: generative AI's largest measured benefits accrue to relatively lower-skilled or lower-performing decision-makers, a pattern Brynjolfsson, Li, and Raymond attribute specifically to the technology's capacity to accelerate the transfer of organizational best practices that experienced decision-makers have already internalized [3], [11]. Table 2's macroeconomic evidence indicates this individual-level pattern scales to a substantial aggregate economic effect concentrated in language-intensive business functions, a finding consistent with Agrawal, Gans, and Goldfarb's characterization of generative AI as fundamentally

a prediction and synthesis technology rather than a general judgment-replacement technology [5], [7], [12]. Table 3's risk-oriented synthesis, however, indicates that this generally positive productivity picture is bounded by the capability-frontier effect Dell'Acqua et al. documented directly: the same generative AI tool that improves performance on well-matched tasks can degrade performance on tasks exceeding its effective capability, particularly when decision-makers extend unwarranted trust to confidently presented but substantively incorrect AI output [4]. This capability-frontier risk is compounded by the organizational governance challenges Berente, Gu, Recker, and Santhanam, and Floridi et al. identify, since generative AI systems' internal reasoning processes remain substantially less auditable than a human decision-maker's stated rationale, a limitation with particular consequence for business decisions carrying legal or fiduciary accountability [10], [15]. Davenport and Ronanki's finding that successful AI implementation depends on deliberate workflow redesign rather than simple tool insertion suggests that Table 3's risk matrix should inform not merely whether an organization adopts generative AI for a given decision task, but how the surrounding decision workflow, verification steps, escalation procedures, and accountability assignment, is structured around that adoption [13].

5. Conclusion

This paper has reviewed the theoretical and empirical literature on generative artificial intelligence's role in modern business decision-making, synthesizing bounded-rationality and prediction-machine theory with field-experimental productivity evidence, macroeconomic potential estimates, and capability-frontier risk findings. The evidence indicates that generative AI's current business value is genuine, substantial, and reproducible across customer operations, consulting, and software engineering functions, but is neither uniform across decision-maker skill levels nor unconditional across task types. The strongest productivity gains are concentrated among relatively lower-skilled or lower-performing decision-makers, a leveling effect with meaningful implications for how organizations think about workforce development and task allocation. The capability-frontier evidence further establishes that generative AI assistance can degrade, rather than improve, decision quality when applied to tasks exceeding the technology's effective competence, particularly given decision-makers' documented tendency to extend unwarranted trust to confidently presented AI output. Collectively, this evidence supports characterizing generative AI's current role in business decision-making as augmentation bounded by a task-specific capability frontier, rather than either a general decision-automation technology or a uniformly beneficial productivity tool.

6. Future Work

Future research should prioritize systematic, task-level mapping of generative AI's capability frontier across specific business decision domains, extending Dell'Acqua et al.'s consulting-specific finding toward the customer operations, marketing, product development, and financial decision-making functions that Chui et al. and Eloundou et al. identify as high-exposure but that have not yet been subjected to comparable controlled capability-frontier testing. Further work is needed to develop and validate organizational workflow-design frameworks that operationalize Davenport and Ronanki's finding that successful AI implementation depends on deliberate process redesign, translating this finding into concrete guidance for the verification steps, escalation procedures, and accountability structures that should surround generative AI-assisted decisions of varying stakes. Continued research should also examine the false-confidence and automation-bias risks Dell'Acqua et al. identified through the lens of decision-maker training interventions, to establish whether targeted training can reduce decision-makers' tendency to accept incorrect AI output without adequate independent verification. Finally, future work should extend Floridi et al.'s explicability principle into concrete, auditable governance standards specifically for generative AI-assisted business decisions carrying legal or fiduciary accountability, an area where current organizational practice appears to lag substantially behind both the technology's capability and its documented risks.

References

1. H. A. Simon, *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization*, 4th ed. New York, NY, USA: Free Press, 1997.
2. D. Kahneman, *Thinking, Fast and Slow*. New York, NY, USA: Farrar, Straus and Giroux, 2011.
3. E. Brynjolfsson, D. Li, and L. R. Raymond, "Generative AI at work," National Bureau of Economic Research Working Paper, no. 31161, 2023.
4. F. Dell'Acqua, E. McFowland III, E. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Kraymer, F. Candelon, and K. R. Lakhani, "Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality," Harvard Business School Working Paper, no. 24-013, 2023.
5. M. Chui, E. Hazan, R. Roberts, A. Singla, K. Smaje, A. Sukharevsky, L. Yee, and R. Zimmel, "The economic potential of generative AI: The next productivity frontier," McKinsey & Company, McKinsey Global Institute report, 2023.

6. S. Peng, E. Kalliamvakou, P. Cihon, and M. Demirer, "The impact of AI on developer productivity: Evidence from GitHub Copilot," arXiv preprint arXiv:2302.06590, 2023.
7. T. Eloundou, S. Manning, P. Mishkin, and D. Rock, "GPTs are GPTs: An early look at the labor market impact potential of large language models," arXiv preprint arXiv:2303.10130, 2023.
8. Y. R. Shrestha, S. M. Ben-Menahem, and G. von Krogh, "Organizational decision-making structures in the age of artificial intelligence," *California Management Review*, vol. 61, no. 4, pp. 66–83, 2019.
9. G. von Krogh, "Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing," *Academy of Management Discoveries*, vol. 4, no. 4, pp. 404–409, 2018.
10. N. Berente, B. Gu, J. Recker, and R. Santhanam, "Managing artificial intelligence," *MIS Quarterly*, vol. 45, no. 3, pp. 1433–1450, 2021.
11. S. Noy and W. Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, vol. 381, no. 6654, pp. 187–192, 2023.
12. A. Agrawal, J. Gans, and A. Goldfarb, "How AI will change the way we make decisions," *Harvard Business Review*, 2018.
13. T. H. Davenport and D. D. D'Ignazio, "Artificial intelligence for the real world," *Harvard Business Review*, vol. 96, no. 1, pp. 108–116, 2018.
14. J. Bughin, E. Hazan, S. Ramaswamy, M. Chui, T. Allas, P. Dahlström, N. Henke, and M. Trench, "Notes from the AI frontier: Insights from hundreds of use cases," McKinsey & Company, McKinsey Global Institute report, 2018.
15. L. Floridi, J. Cowls, M. Beltrametti, R. Chatila, P. Chazerand, V. Dignum, C. Luetge, R. Madelin, U. Pagallo, F. Rossi, B. Schafer, P. Valcke, and E. Vayena, "AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations," *Minds and Machines*, vol. 28, no. 4, pp. 689–707, 2018.
16. A. Susarla, R. Gopal, J. B. Thatcher, and S. Sarker, "The Janus effect of generative AI: Charting the path for responsible conduct of scholarly research in information systems," *Information Systems Research*, vol. 34, no. 2, pp. 399–408, 2023.
17. H. S. Sætra, "Generative AI: Here to stay, but for good?," *Technology in Society*, vol. 75, 102372, 2023.