

REVIEW

Open access

Generative AI in Digital Business and Management Studies: A Systematic Review of Productivity, Decision Quality, Governance, and Organizational Risk

Charlotte Dupuis^{1*}, Hugo Perrin¹

Abstract

Generative artificial intelligence has moved rapidly from experimental use to practical adoption across digital business and management contexts. Its diffusion has been accelerated by large language models, image generators, code generators, and conversational systems that can support content creation, analysis, automation, and decision support. This systematic review examines the evidence on Generative AI in business and management studies, with particular attention to productivity, decision quality, governance, and organizational risk. The review addresses the need for a balanced synthesis that recognises both the performance promise of Generative AI and the risks created by its probabilistic, opaque, and adaptive nature. The findings show that Generative AI can improve productivity by reducing task completion time, expanding output volume, supporting creative work, and assisting knowledge workers. However, the evidence also indicates uneven benefits across tasks, expertise levels, organizational contexts, and governance conditions, while decision quality remains vulnerable to hallucination, bias, over-reliance, and weak accountability. The review concludes that Generative AI should be understood not merely as a productivity technology but as an organizational transformation phenomenon. Its business value depends on the co-development of human oversight, governance structures, risk controls, workforce capabilities, and context-sensitive implementation practices.

Keywords AI governance, Productivity, Generative AI, Large language models, Decision quality, Organizational risk

*Correspondence:

Charlotte Dupuis
charlotte.dupuis@univ-lille.fr

¹ Department of Digital Strategy and Business Innovation, Faculty of Economics, University of Lille, Lille, France

Introduction

Generative AI has become a defining technology in digital business because it can produce text, images, code, recommendations, summaries, and analytical outputs rather than merely classify or predict existing data. Feuerriegel, Hartmann, Janiesch, and Zschech conceptualise Generative AI as a broad class of systems capable of generating novel digital artefacts, while Teubner, Flath, Weinhardt, van der Aalst, and Hinz argue that ChatGPT marks a shift in how organizations interact with intelligent systems [1, 2]. This development has moved AI from specialised technical domains into everyday managerial and knowledge-work settings. As a result,

Generative AI has become relevant to strategy, operations, marketing, human resources, innovation, governance, and organizational design.

The business literature has responded quickly, but the evidence remains fragmented across disciplinary silos. Marketing studies examine customer engagement, content generation, and consumer behaviour, while management and entrepreneurship research focus on innovation, strategic decision-making, and venture communication [3–6]. Human resource management scholarship has begun to analyse how Generative AI alters recruitment, work design, skills, and employment relations [7, 8]. Information systems

and governance research, meanwhile, highlights the need to manage algorithmic opacity, organizational accountability, and responsible AI implementation [9–11].

The objective of this systematic review is to synthesise peer-reviewed evidence on Generative AI in digital business and management studies, focusing on four interrelated themes: productivity, decision quality, governance, and organizational risk. Productivity is treated broadly, including time savings, output quality, task automation, augmentation, and changes in knowledge-work performance, as illustrated by experimental and field evidence on AI-assisted work [12, 13]. Decision quality includes accuracy, creativity, bias, judgment, and human-AI oversight, building on work that examines AI's role in strategic evaluation and organizational decision-making [14–16]. Governance and risk include hallucination, misinformation, accountability, intellectual property, data provenance, workforce disruption, reputational exposure, and ethical dilemmas [9, 17, 18].

This review is guided by four research questions. First, how is Generative AI being applied across digital business and management functions? Second, what evidence exists regarding its effects on productivity and decision quality? Third, what governance challenges and organizational risks are identified in the literature? Fourth, what theoretical, empirical, and managerial gaps remain for future research? The article proceeds by describing the review method, mapping the application landscape, synthesising productivity and decision-quality evidence, analysing governance and risk, consolidating research gaps, developing a future research agenda, and translating the findings into managerial implications.

Review Method and Selection Criteria

This systematic review followed a structured review logic designed to identify, screen, appraise, and synthesise peer-reviewed journal evidence on Generative AI in digital business and management studies. The review design was informed by systematic literature review guidance that emphasises transparent search strategy, explicit inclusion criteria, quality appraisal, and synthesis logic [19]. It also drew on the SPAR-4-SLR protocol, which supports assembling, arranging, and assessing literature in a reproducible manner [20]. The aim was not to produce new empirical findings but to organise existing evidence into a

coherent synthesis for business and management scholarship.

The search strategy targeted peer-reviewed journal articles published between 2017 and 2026 that addressed Generative AI, large language models, ChatGPT, AI-assisted work, AI governance, organizational risk, productivity, decision-making, and digital business applications. The search logic prioritised journals in information systems, management, strategy, marketing, organization studies, human resource management, innovation, entrepreneurship, AI governance, and business research. Articles were included when they directly addressed Generative AI or provided essential conceptual foundations for understanding AI in organizations, algorithmic work, human-AI collaboration, or responsible AI governance [11, 15, 21, 22]. Articles were excluded when they were books, policy reports, websites, theses, non-peer-reviewed manuscripts, or purely technical studies without business or management relevance.

The screening process was organised around relevance to the review's four focal constructs: productivity, decision quality, governance, and organizational risk. Empirical articles were appraised for setting, task type, evidence strength, outcome measures, and generalisability, while conceptual articles were assessed for theoretical contribution, relevance to Generative AI, and usefulness for business and management synthesis [12–14]. Governance-oriented articles were retained when they addressed accountability, auditing, policy, risk, responsible AI, or organizational control mechanisms [9, 10, 17]. Broader AI management articles were included when they provided conceptual foundations for automation, augmentation, algorithmic control, human-AI symbiosis, or organizational theorising [15, 16, 23].

Table 1 summarises the systematic review search strategy, inclusion criteria, and final sample. The final sample contained exactly 36 peer-reviewed journal articles, selected to balance direct Generative AI studies with foundational management, information systems, governance, and review-method contributions. The evidence base covers empirical productivity studies, conceptual management frameworks, marketing and innovation applications, entrepreneurship use cases, human resource management implications, governance frameworks, and systematic review methodology [1, 3, 7, 9, 12, 19, 20, 24]. This selection enables a systematic synthesis of Generative AI as both a digital business

opportunity and an organizational risk-management challenge.

Table 1. Systematic Review Search Strategy and Selection Criteria: Databases, Search Terms, Screening Process, and Final Sample Composition

Review element	Specification used in this systematic review	Rationale for inclusion in the review design	Final treatment in the synthesis
Time window	2017–2026	Captures foundational AI-in-organization studies and the rapid post-2022 emergence of Generative AI research	Articles were included only if published within the specified window
Literature type	Peer-reviewed journal articles only	Ensures scholarly quality and excludes non-reviewed reports, websites, theses, and informal commentary	Final sample restricted to 36 journal articles
Core search concepts	Generative AI, ChatGPT, large language models, productivity, decision quality, governance, organizational risk, systematic review	Aligns the search with the article's four substantive review themes	Search terms were combined across technology, outcome, governance and method domains
Business domains	Management, information systems,	Reflects Generative AI's cross-	Articles were grouped by application

	strategy, marketing, human resources, innovation, entrepreneurship, organization studies	functional diffusion in digital business	domain and evidence contribution
Inclusion criteria	Direct relevance to Generative AI or strong conceptual relevance to AI in organizations, governance, decision-making, or systematic review method	Allows the review to combine emerging Generative AI evidence with necessary theoretical foundations	Direct Generative AI studies and foundational AI management studies were both retained
Exclusion criteria	Books, theses, policy reports, websites, arXiv papers, conference papers, purely technical AI studies, and articles without business relevance	Maintains the review's focus on peer-reviewed business and management evidence	Non-journal and non-business sources were excluded
Screening focus	Productivity, decision quality, governance, organizational risk, application domain, and methodological relevance	Ensures that all retained articles contribute to at least one review question	Articles were coded according to their primary and secondary contribution
Quality appraisal	Relevance, evidence type, theoretical contribution, empirical setting, outcome clarity, and governance relevance	Supports critical synthesis rather than descriptive listing	Evidence was synthesised narratively and thematically
Final sample	36 peer-reviewed journal articles	Satisfies the reference	All 36 articles are

		requirement while covering the review's full conceptual scope	used across the complete manuscript synthesis
--	--	---	---

Generative AI in Digital Business Context

The reviewed literature shows that Generative AI has entered digital business as a general-purpose capability rather than a single application. Feuerriegel, Hartmann, Janiesch, and Zschech describe Generative AI as a technology class that can generate digital content across text, images, software, and synthetic outputs, while Teubner, Flath, Weinhardt, van der Aalst, and Hinz emphasise the managerial significance of conversational systems such as ChatGPT [1, 2]. In business contexts, this means that Generative AI can support both front-office and back-office activities. Its relevance therefore extends from customer communication and marketing content to coding, analytics, knowledge management, and managerial decision support.

Marketing and customer-facing domains are among the most visible areas of adoption. Hermann and Puntoni argue that marketing research is shifting from predictive AI toward generative forms of customer interaction, while Grewal, Saturnino, Davenport, and Guha examine how Generative AI may reshape marketing practice and consumer engagement [3, 25]. Cillo and Rubera identify research opportunities around Generative AI in innovation and marketing processes, and Heitmann, Jansen, Reisenbichler, and Schweidel show how visual Generative AI can be used to engage customers through image-based content [4, 26]. These studies suggest that business value is not limited to efficiency but also includes creativity, personalization, and experimentation.

Generative AI is also becoming relevant to innovation, entrepreneurship, and strategic management. Mariani and Dwivedi position Generative AI as a catalyst for innovation management research, while Ferrati, Kim, and Muffatto discuss its role in intelligence augmentation for entrepreneurship research [5, 24]. Etemad examines whether international entrepreneurial enterprises should adopt Generative AI, and Short and Short analyse how ChatGPT and prompt engineering can support

entrepreneurial rhetoric creation [6, 27]. In strategy, Doshi, Bell, Mirzayev, and Vanneste show that Generative AI can be used to evaluate strategic decisions, indicating that the technology is moving from content generation toward higher-level managerial judgment [14].

Table 2 maps the application domains of Generative AI in business and management. Across the reviewed evidence, the main application domains include marketing, innovation, entrepreneurship, human resources, knowledge work, decision support, governance, and organizational control [6-8, 24, 25]. However, the application landscape is uneven because marketing and productivity studies are more developed than longitudinal studies of organizational transformation, risk governance, and workforce redesign [12, 28, 29]. This imbalance indicates that Generative AI research has moved quickly in identifying use cases but more slowly in explaining how these applications become embedded in organizational routines, capabilities, and governance systems.

Table 2. Generative AI Applications in Digital Business: Domains, Use Cases, and Underlying Technologies

Business or management domain	Representative Generative AI use cases	Underlying technologies	Management implications
Marketing and customer engagement	Automated content creation, personalised campaigns, synthetic images, customer-facing communication	Large language models, visual Generative AI, multimodal systems	Fast personalization, cost reduction, enhanced engagement
Innovation management	Ideation, concept development, market scanning, product design support, innovation process augmentation	Large language models, image generation, text-to-concept systems	Expanded creative capacity, faster experimentation, intelligent augmentation

Entrepreneurship	Pitch development, venture communication, opportunity framing, entrepreneurial rhetoric	ChatGPT, prompt engineering, conversational AI	Low
Strategic decision support	Scenario evaluation, strategic alternative assessment, decision explanation, managerial analysis	Large language models, decision-support interfaces	Fast and general
Human resource management	Job description drafting, employee communication, skills analysis, recruitment support, workforce planning	Large language models, conversational AI, text generation systems	Administrative efficiency and augmentation
Knowledge work and professional services	Writing assistance, summarisation, coding, document drafting, analytical support	Large language models, code generation, text generation systems	Reduced completion time, improved volume, faster experimentation
Governance and compliance	AI auditing, policy drafting, risk classification, accountability support, governance reporting	Large language models, audit frameworks, responsible AI tools	Improved visibility and documentation
Organizational control and work design	Algorithmic supervision, coordination, workflow automation,	Learning algorithms, Generative AI interfaces,	Efficient coordination and scaling

	human-AI collaboration	workplace AI systems	
--	------------------------	----------------------	--

Productivity and Decision Quality Outcomes

The strongest evidence on productivity indicates that Generative AI can improve task speed and output quality, particularly in structured writing, customer service, and professional knowledge-work settings. Brynjolfsson, Li, and Raymond show that Generative AI assistance can raise productivity in customer-support work, especially for less experienced workers, suggesting that the technology may compress learning curves rather than simply benefit already skilled employees [12]. Noy and Zhang similarly demonstrate that Generative AI can reduce completion time and improve writing quality in professional writing tasks [13]. These findings support the view that Generative AI operates as an augmentation technology when tasks are bounded, feedback is available, and outputs can be evaluated by humans.

The productivity evidence is not uniformly positive because performance gains depend on task structure, worker expertise, and the fit between AI capability and task requirements. Dell’Acqua, McFowland, Mollick, Lifshitz-Assaf, Kellogg, Rajendran, Kraye, Candelon, and Lakhani describe the “jagged technological frontier,” where AI improves performance on some tasks but can reduce quality when users apply it beyond its reliable competence boundary [30]. Przegalinska, Triantoro, Kovbasiuk, Ciechanowski, Freeman, and Sowa also frame collaborative AI through resource-based and task-technology fit perspectives, showing that performance depends on alignment between organizational resources, task demands, and AI use [28]. This suggests that Generative AI productivity should be treated as conditional rather than automatic.

Decision quality is similarly mixed because Generative AI can broaden the range of alternatives considered while also generating persuasive but unreliable outputs. Doshi, Bell, Mirzayev, and Vanneste show that Generative AI can assist in evaluating strategic decisions, but this use requires careful attention to judgment, framing, and human evaluation [14]. Raisch and Krakowski’s automation–augmentation paradox remains relevant because managerial decision quality depends on whether AI

replaces human judgment or strengthens it through structured collaboration [15]. Jarrahi’s account of human-AI symbiosis further implies that decision benefits arise when human contextual reasoning and machine-generated analysis are combined rather than treated as substitutes [16].

Table 3 synthesises the evidence on productivity and decision quality outcomes of Generative AI. The reviewed studies show that positive outcomes are most credible when the task is well specified, the user has enough expertise to evaluate output quality, and the organization has oversight mechanisms that prevent blind reliance [12, 13, 28, 30]. Conversely, decision quality can deteriorate when users defer to confident AI outputs, apply models outside their competence boundary, or fail to detect fabricated reasoning [14, 18, 30]. The evidence therefore supports a balanced interpretation: Generative AI can enhance productivity and decision support, but only under conditions of task fit, user capability, and disciplined governance.

Table 3. Productivity and Decision Quality Outcomes of Generative AI: Empirical Evidence, Effect Sizes, and Moderating Factors

Outcome area	Evidence pattern in the reviewed literature	Indicative effect or contribution	Moderating factors
Task completion time	Generative AI reduces time required for writing, summarisation, and customer-support tasks	Strongest evidence appears in controlled and field studies of bounded knowledge work	Task complexity, prompt quality, familiarity, availability, review
Output quality	AI assistance can improve average output quality, especially for less experienced users	Quality gains are visible in writing and customer-service evidence	Baseline worker evaluation, standard feedback loop

Knowledge-work augmentation	Generative AI supports drafting, ideation, coding, synthesis, and analytical preparation	Evidence suggests augmentation rather than full substitution in complex work	Human expert task completion, organizational norms
Strategic decision support	AI can generate alternatives, organise reasoning, and assist evaluation of strategic options	Evidence remains early and context-sensitive	Management judgment, decision stakes, domain knowledge, model reliability
Creativity and ideation	AI supports idea generation, entrepreneurial rhetoric, and innovation process exploration	Conceptual and applied evidence is stronger than longitudinal evidence	Novel criteria, human selection, market knowledge
Bias and error	AI may reduce some human bottlenecks but can also reproduce bias or fabricate information	Risk is especially salient in decisions requiring accuracy, fairness, or accountability	Training prompts, framing, oversight, auditability
Expertise effects	Less experienced workers may benefit disproportionately in bounded settings	Evidence suggests productivity equalisation in some work contexts	Prior expertise, task evaluation, coaching, feedback
Reliability boundary	AI performance is uneven across tasks and can decline when used beyond its capability frontier	The “jagged frontier” is central to interpreting productivity evidence	Task type uncertainty, ambiguous evaluation, difficulty

Governance and Organizational Risk

The governance literature emphasises that Generative AI creates distinctive risks because it generates plausible outputs without guaranteeing truth, provenance, accountability, or legal clarity. Taeihagh frames governance of Generative AI as a policy and institutional challenge, while Janssen conceptualises responsible governance by treating GenAI as a complex adaptive system rather than a controllable standalone tool [9, 10]. Floridi and Chiriatti’s analysis of GPT-3 remains foundational because it clarifies the limits, scope, and consequences of large generative models [18]. These studies indicate that governance must address uncertainty, not merely compliance.

A central organizational risk is that Generative AI systems can produce hallucinated, biased, or untraceable content that enters business decisions, customer communications, or public-facing materials. Mökander, Schuett, Kirk, and Floridi propose a three-layered approach to auditing large language models, highlighting the need to examine governance at model, application, and organizational levels [17]. Berente, Gu, Recker, and Santhanam argue that managing AI requires attention to organizational structures, routines, and accountability, not just technical accuracy [11]. In management settings, this means that governance must define who is responsible for AI-generated outputs, how outputs are reviewed, and when human approval is mandatory.

Organizational risk also includes workforce displacement, algorithmic control, skill erosion, and contested authority. Budhwar, Chowdhury, Wood, Aguinis, Bamber, Beltran, Boselie, Cooke, Decker, DeNisi, Dey, Guest, Knoblich, Malik, Paauwe, Papagiannidis, Patel, Pereira, Ren, Rogelberg, Saunders, Tung, and Varma show that Generative AI raises major human resource management questions around work design, employment relations, skills, and fairness [8]. Kellogg, Valentine, and Christin’s analysis of algorithms at work helps explain why AI-enabled systems may become a new terrain of control and resistance [21]. Faraj, Pachidi, and Sayegh similarly show that learning algorithms reshape organizing by changing how work is coordinated, evaluated, and contested [22].

Table 4 catalogues governance challenges and organizational risks associated with Generative AI. The reviewed evidence suggests that risk mitigation requires layered controls, including model evaluation, human review,

documentation, audit trails, user training, access controls, escalation rules, and clear accountability structures [9-11, 17]. Khanal, Zhang, and Taeihagh add that Big Tech power in the policy process complicates governance because key capabilities, infrastructures, and standards are shaped by dominant technology providers [29]. Therefore, organizational governance must be both internal and ecosystem-aware.

Table 4. Governance Challenges and Organizational Risks of Generative AI: Types, Severity, and Mitigation Approaches

Governance or risk category	Description in business and management contexts	Potential severity	Main affected stakeholders
Hallucination and factual unreliability	AI generates plausible but false or unsupported outputs	High in decision, legal, financial, and customer-facing settings	Management, customer, regulatory, employee
Data provenance and traceability	Generated outputs may lack clear origin, evidence base, or documentation	High where auditability is required	Compliance teams, regulatory, customer
Intellectual property uncertainty	Generated text, images, or code may raise ownership or infringement concerns	Medium to high depending on use case	Legal teams, creator, firms, client
Accountability ambiguity	Responsibility for AI-generated outputs may be unclear	High in strategic, HR, and public-facing decisions	Senior management, employee, customer
Bias and unfair treatment	AI may reproduce or amplify social	High in HR, customer segmentation,	Employee, applicant, customer

	organizational, or data-driven bias	and service allocation	
Reputational damage	Poor AI-generated content may be misleading, offensive, or brand-inconsistent	High in marketing and communication	Customer investor public audience
Workforce displacement and deskilling	AI may automate tasks, alter roles, or weaken professional expertise	Medium to high across knowledge-work functions	Employee management HR leadership
Over-reliance and automation bias	Users may defer to AI even when outputs are weak or fabricated	High in complex decision contexts	Management professional customer
Security and misuse	AI tools may expose sensitive data or generate harmful content	High in regulated and confidential domains	Firms customer regulatory
Ecosystem dependence	Firms may depend on external providers for models, infrastructure, and standards	Medium to high in platform-dependent environments	Firms platform regulatory

Paliszkiewicz, Wach, and Ziemba explicitly examine Generative AI as a new context for management theories [1, 31]. However, many studies still treat Generative AI as a tool rather than as an organizational phenomenon that changes routines, authority, capability development, and governance. This creates a gap between technological novelty and management theory-building.

The evidence base also shows a methodological imbalance because short-term experiments and conceptual essays are more common than longitudinal organizational studies. Brynjolfsson, Li, and Raymond and Noy and Zhang provide important evidence on productivity effects, but their findings need to be extended across longer time horizons, diverse industries, and different organizational maturity levels [6, 8]. Dell'Acqua, McFowland, Mollick, Lifshitz-Assaf, Kellogg, Rajendran, Kraye, Candelon, and Lakhani show that outcomes vary across the jagged frontier of AI capability, making longitudinal and task-level research especially important [7]. Without such studies, the literature may overestimate immediate gains and underestimate adaptation costs.

Another gap concerns the relationship between Generative AI and decision quality in high-stakes managerial contexts. Doshi, Bell, Mirzayev, and Vanneste examine AI-assisted evaluation of strategic decisions, but the field still lacks comparative studies of when Generative AI improves judgment, when it introduces bias, and when it encourages overconfidence [11]. Lebovitz, Levina, and Lifshitz-Assaf show that AI systems may be evaluated against unstable or misleading versions of expert ground truth, which complicates claims about decision accuracy [32]. This problem is especially relevant to Generative AI because its outputs can appear analytically persuasive even when the underlying reasoning is unreliable.

Table 5 consolidates the identified research gaps and maps them to future research directions. The largest gaps concern longitudinal impact, organizational implementation, human-AI role redesign, governance effectiveness, ecosystem dependence, and the co-evolution of AI capabilities and organizational routines [10, 23, 29, 30]. Von Krogh's call for phenomenon-based theorising in AI research remains particularly relevant because Generative AI is not only a technical artifact but a changing organizational phenomenon that requires new explanatory concepts [23]. Future research should therefore move from documenting use cases toward explaining causal

Synthesis of Evidence and Research Gaps

The synthesis reveals a clear imbalance between the rapid expansion of application-oriented studies and the slower development of cumulative theory. Feuerriegel, Hartmann, Janiesch, and Zschech provide a broad conceptual foundation for understanding Generative AI, while Korzynski, Mazurek, Altmann, Ejdays, Kazlauskaitė,

mechanisms, boundary conditions, and governance outcomes.

Table 5. Synthesis of Research Gaps in Generative AI Business Research: Thematic Deficits and Recommended Future Research Directions

Research gap	Current limitation in the reviewed literature	Why the gap matters	Recommended future directions
Limited longitudinal evidence	Many studies examine early adoption, experiments, or conceptual implications	Short-term gains may not predict long-term capability, risk, or workforce effects	Conduct year-long studies on Generative AI adoption in institutional settings
Weak organizational implementation theory	Use cases are often described without explaining how organizations embed AI into routines	Implementation determines whether AI becomes productive, resisted, or risky	Study the change management, governance, and structural role of AI in organizations
Underdeveloped decision-quality evidence	Evidence remains limited on strategic, high-stakes, and ambiguous decisions	AI may improve speed while weakening judgment or accountability	Conduct experiments on decision assistance only, and develop decision process control mechanisms
Insufficient governance evaluation	Governance frameworks are proposed more often than empirically tested	Firms need evidence on which controls actually reduce risk	Test governance models with workflow provisions, system oversight, and privacy protection
Limited workforce transformation research	HR studies identify implications but lack detailed	Generative AI may alter skills, authority, identity, and	Explore the reskilling and deskilling of the workforce

	evidence on long-term work redesign	employment relations	em part
Narrow productivity measurement	Productivity is often measured through time and output quality	Business value also includes learning, creativity, risk, coordination, and customer outcomes	De multidisciplinary productivity measurement and knowledge
Weak attention to ecosystem dependence	Firms rely on model providers, cloud platforms, and external standards	Provider power affects governance, risk, cost, and strategic control	Study dependence on platform governance and AI ecosystem
Limited sectoral and contextual comparison	Evidence is concentrated in general knowledge work and selected business functions	AI value and risk differ across industries, cultures, and regulation levels	Conduct comparative studies across sectors, sizes, and regulatory environments
Insufficient theory of human-AI collaboration	Augmentation is widely discussed but weakly specified	Collaboration quality determines whether AI improves or harms outcomes	Build theories on the role and external interaction calibration of AI
Underexplored ethical and reputational risk	Risk is recognised but not always connected to measurable organizational consequences	Reputational failures can undermine adoption and stakeholder trust	Link failures to customer brand and employee responsibility

Figure 1 integrates the review's evidence on Generative AI applications, productivity and decision quality outcomes, governance requirements, organizational risks, and unresolved research gaps into a unified value–risk framework.




Figure 1. Generative AI Value–Risk Governance Framework for Digital Business and Management

Future Research Agenda

Future research should prioritise longitudinal studies that track Generative AI from pilot adoption to organizational scaling. Current productivity evidence from bounded tasks is valuable, but it does not fully explain how benefits evolve after workers adapt, managers redesign processes, and organizations institutionalise AI use [12, 13, 30]. Longitudinal research should examine whether early productivity gains persist, decline, or transform into new forms of coordination cost. It should also assess whether Generative AI produces durable capabilities or merely short-term efficiency effects.

A second research direction concerns multi-level theorising across individuals, teams, organizations, and ecosystems. Raisch and Krakowski's automation–augmentation paradox provides a strong foundation for understanding competing logics of AI use, but Generative AI requires more precise theories of how augmentation operates across managerial layers [15]. Berente, Gu, Recker, and Santhanam show that AI management requires organizational attention, while von Krogh argues that AI provides opportunities for phenomenon-based theorising [11, 23]. Future research should therefore explain how individual prompting practices, team review routines, organizational governance, and platform ecosystems interact.

A third priority is comparative governance research. Taeihagh and Janssen provide important governance frameworks, but future studies should compare how different firms design, implement, and evaluate governance models for Generative AI [9, 10]. Mökander, Schuett, Kirk, and Floridi's auditing framework could be extended into empirical studies that test how audit layers work in real organizational settings [17]. Researchers should examine whether governance effectiveness varies by sector, task criticality, regulation, model architecture, and data sensitivity.

A fourth research stream should study the co-evolution of Generative AI capabilities and organizational design. Chowdhury, Budhwar, and Wood highlight the strategic human resource management implications of Generative AI, and Budhwar and colleagues show that HRM must address new questions around skills, roles, and employment relations [7, 8]. Kellogg, Valentine, and Christin's work on algorithmic control suggests that AI adoption may also reshape authority, surveillance, and resistance at work [21]. Future research should examine how firms redesign roles, decision rights, training systems, and accountability structures as Generative AI becomes embedded in everyday work.

Managerial Implications

Managers should approach Generative AI as a strategic capability rather than an isolated productivity tool. The reviewed evidence shows that AI can improve speed and output quality, but only when use cases are matched to task characteristics and embedded in appropriate review systems [12, 13, 28, 30]. A practical Generative AI strategy should classify tasks according to risk, evaluability, business value, and required human expertise. This prevents firms from applying AI indiscriminately across functions where reliability, accountability, or contextual judgment may be weak.

Managers also need governance frameworks that address hallucination, provenance, accountability, and reputational exposure before widespread deployment. Governance should include approved tools, data-use rules, output verification, audit trails, escalation pathways, and named human responsibility for AI-supported decisions [9, 11, 17]. Floridi and Chiriatti's analysis of GPT-3's limits reinforces the need to treat model outputs as probabilistic artefacts rather than authoritative knowledge [18]. In practice, this means that AI-generated material should be reviewed differently depending on whether it is used for internal brainstorming, customer communication, compliance work, or strategic decisions.

Workforce capability development is another managerial priority. Generative AI can support employees by reducing routine effort and expanding access to analytical and creative support, but it can also create anxiety, deskilling, and contested control if implemented without participation [8, 21, 22]. Human resource leaders should therefore combine reskilling, role redesign, employee consultation,

and clear communication about augmentation versus substitution. This is especially important because Generative AI affects not only tasks but also professional identity, autonomy, and trust.

Finally, managers should pilot Generative AI before scaling it and maintain human oversight for critical decisions. Evidence on the jagged technological frontier suggests that managers must learn where AI performs reliably and where it produces hidden errors [30]. Strategic, legal, financial, HR, and public-facing uses should require stronger validation than low-risk ideation or drafting tasks [14, 17]. A disciplined pilot-to-scale approach allows firms to capture productivity gains while building the governance maturity needed to manage organizational risk.

Conclusion

This systematic review shows that Generative AI has substantial potential to reshape digital business and management. The evidence indicates meaningful opportunities for productivity improvement, knowledge-work augmentation, creative support, marketing innovation, entrepreneurial communication, and managerial decision support. At the same time, these benefits are conditional on task fit, user expertise, human oversight, and organizational readiness.

The review also shows that governance and risk management have not developed as quickly as adoption. Hallucination, data provenance, intellectual property uncertainty, accountability ambiguity, workforce disruption, reputational exposure, and over-reliance remain central concerns. Generative AI should therefore be understood as

a powerful but risky organizational technology whose value depends on the co-evolution of adoption, governance, skills, and accountability.

Future research should move beyond early enthusiasm and isolated use cases toward longitudinal, comparative, and theory-driven studies of Generative AI in organizations. Managers should proceed with ambition but not with blind optimism, building governance frameworks and workforce capabilities alongside experimentation. The central lesson of this review is that Generative AI's business value will be realised only when productivity, decision quality, governance, and organizational risk are managed together.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 01 Nov 2025 Revised: 10 Dec 2025 Accepted: 20 Jan 2026

Published online: 18 March 2026

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Feuerriegel S, Hartmann J, Janiesch C, Zschech P. Generative AI. *Bus Inf Syst Eng.* 2024;66(1):111-26.

Teubner T, Flath CM, Weinhardt C, Van Der Aalst W, Hinz O. Welcome to the Era of ChatGPT et al. *Bus Inf Syst Eng.* 2023;65(2):95-101.

- Hermann E, Puntoni S. Artificial intelligence and consumer behavior: From predictive to generative AI. *J Bus Res.* 2024;180:114720.
- Cillo P, Rubera G. Generative AI in innovation and marketing processes: A roadmap of research opportunities. *J Acad Mark Sci.* 2025;53(3):684-701.
- Ferrati F, Kim PH, Muffatto M. Generative AI in entrepreneurship research: Principles and practical guidance for intelligence augmentation. *Found Trends Entrep.* 2024;20(3):245-383.
- Short CE, Short JC. The artificially intelligent entrepreneur: ChatGPT, prompt engineering, and entrepreneurial rhetoric creation. *J Bus Ventur Insights.* 2023;19:e00388.
- Chowdhury S, Budhwar P, Wood G. Generative artificial intelligence in business: Towards a strategic human resource management framework. *Br J Manag.* 2024;35(4):1680-91.
- Budhwar P, Chowdhury S, Wood G, Aguinis H, Bamber GJ, Beltran JR, et al. Human resource management in the age of generative artificial intelligence: Perspectives and research directions on ChatGPT. *Hum Res Manag J.* 2023;33(3):606-59.
- Taeihagh A. Governance of generative AI. *Policy Soc.* 2025;44(1):1-22.
- Janssen M. Responsible governance of generative AI: conceptualizing GenAI as complex adaptive systems. *Policy Soci.* 2025;44(1):38-51.
- Berente N, Gu B, Recker J, Santhanam R. Managing artificial intelligence. *MIS Q.* 2021;45(3):1433-50.
- Brynjolfsson E, Li D, Raymond L. Generative AI at work. *Q J Econ.* 2025;140(2):889-942.
- Noy S, Zhang W. Experimental evidence on the productivity effects of generative artificial intelligence. *Science.* 2023;381(6654):187-92.
- Doshi AR, Bell JJ, Mirzayev E, Vanneste BS. Generative artificial intelligence and evaluating strategic decisions. *Strateg Manag J.* 2025;46(3):583-610.
- Raisch S, Krakowski S. Artificial intelligence and management: The automation-augmentation paradox. *Acad Manag Rev.* 2021;46(1):192-210.
- Jarrahi MH. Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Bus Horiz.* 2018;61(4):577-86.
- Mökander J, Schuett J, Kirk HR, Floridi L. Auditing large language models: A three-layered approach. *AI Ethics.* 2024;4(4):1085-115.
- Floridi L, Chiriatti M. GPT-3: Its nature, scope, limits, and consequences. *Mind Mach.* 2020;30(4):681-94.
- Snyder H. Literature review as a research methodology: An overview and guidelines. *J Bus Res.* 2019;104:333-9.
- Paul J, Lim WM, O'Cass A, Hao AW, Bresciani S. Scientific procedures and rationales for systematic literature reviews (SPAR-4-SLR). *Int J Consum Stud.* 2021;45(4):O1-6.
- Kellogg KC, Valentine MA, Christin A. Algorithms at work: The new contested terrain of control. *Acad Manag Annal.* 2020;14(1):366-410.
- Faraj S, Pachidi S, Sayegh K. Working and organizing in the age of the learning algorithm. *Inf Organ.* 2018;28(1):62-70.
- Von Krogh G. Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Acad Manag Discov.* 2018;4(4):404-9.
- Mariani M, Dwivedi YK. Generative artificial intelligence in innovation management: A preview of future research developments. *J Bus Res.* 2024;175:114542.
- Grewal D, Saturnino CB, Davenport T, Guha A. How generative AI is shaping the future of marketing. *J Acad Mark Sci.* 2025;53(3):702-22.
- Heitmann M, Jansen TP, Reisenbichler M, Schweidel DA. EXPRESS: Picture perfect: engaging customers with visual generative AI. *J Mark.* 2025;00222429251356993.
- Etemad H. Transformative potentials of generative artificial intelligence: Should international entrepreneurial enterprises adopt GEN. AI?. *J Int Entrep.* 2024;22(2):141.
- Przegalinska A, Triantoro T, Kovbasiuk A, Ciechanowski L, Freeman RB, Sowa K. Collaborative AI in the workplace: Enhancing organizational performance through resource-based and task-technology fit perspectives. *Int J Inf Manag.* 2025;81:102853.
- Khanal S, Zhang H, Taeihagh A. Why and how is the power of Big Tech increasing in the policy process? The case of generative AI. *Policy Soc.* 2025;44(1):52-69.

Dell'Acqua F, McFowland III E, Mollick E, Lifshitz H, Kellogg KC, Rajendran S, et al. Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organ Sci.* 2026;37(2):403-23.

Korzynski P, Mazurek G, Altmann A, Ejdys J, Kazlauskaite R, Paliszkievicz J, et al. Generative artificial intelligence as a new

context for management theories: Analysis of ChatGPT. *Cent Eur Manag J.* 2023;31(1):3-13.

Lebovitz S, Levina N, Lifshitz-Assaf H. Is AI ground truth really true? The dangers of training and evaluating AI tools based on experts' know-what. *MIS Q.* 2021;45(3):1501-26.