

A Path from Generative AI to Agentic AI in Marketing: A Critical Review, Capability Framework and Future Research Agenda

Ankita Garg

Assistant Professor, MERI College of Engineering and Technology, Haryana, India

Correspondence: ankitagarg644@gmail.com



<https://doi.org/10.55041/ijst.v2i8.036>

Cite this Article: Garg, A. (2026). A Path from Generative AI to Agentic AI in Marketing: A Critical Review, Capability Framework and Future Research Agenda. International Journal of Science, Strategic Management and Technology, 02(8), 1-9. <https://doi.org/10.55041/ijst.v2i8.036>

License:  This article is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited.

Abstract

Purpose. Marketing organizations are moving from prompt-driven Generative Artificial Intelligence (Generative AI) toward agentic systems that plan, act, and adapt across multi-step tasks with limited human supervision. The scholarly literature on this transition is dispersed across marketing, information systems, and computer science, and it rarely distinguishes technological capability from the organizational conditions under which autonomous agents create legitimate and durable marketing value. This article critically reviews that relationship and develops a capability-based framework for autonomous marketing.

Design/methodology/approach. The article is a critical narrative review rather than a systematic or empirical review. Google Scholar was the principal literature-discovery platform, supplemented by backward and forward citation searching, during July–August 2026. The review audits the 35 references in the source manuscript, retains all as traceable and topically relevant, and adds 8 supplementary scholarly sources identified through citation chaining, producing a cited corpus of 43 scholarly works. A structured narrative process covered identification, relevance screening, design-sensitive appraisal, thematic coding, and narrative synthesis. Dynamic capabilities theory structures the analysis.

Findings. Agentic AI does not automatically improve marketing outcomes. Its value is contingent on an organizational capability-conversion process in which AI affordances are translated, through data infrastructure, analytical and creative skill, redesigned workflows, and governance, into accountable marketing routines. Evidence specific to Agentic AI in marketing remains overwhelmingly conceptual, vendor-authored, or transferred from adjacent domains such as personal selling and customer service; controlled, longitudinal, or audited evidence of autonomous-agent performance in marketing is scarce, and widely circulated efficiency and return-on-investment figures originate predominantly from industry and consultancy reports rather than peer-reviewed research.

Originality/value. The review replaces a linear technology-adoption narrative with a capability-conversion explanation, separates reviewed evidence from proposed relationships, and contributes a four-stage evolutionary model, a governance-sensitive AI–capability–marketing relationship framework, and testable propositions distinguishing technological affordances from organizational capability, marketing practice redesign, and verified outcomes.

Keywords

artificial intelligence; generative AI; agentic AI; AI agents; marketing; autonomous marketing; dynamic capabilities; responsible AI

1. Introduction

Artificial intelligence (AI) has progressed from a supporting analytical tool to a technology increasingly proposed as a strategic organizational capability in marketing, shaping customer relationship management, campaign optimization, demand forecasting, and personalized experience (Ma & Sun, 2020; Kaplan & Haenlein, 2019; Davenport et al., 2020). Huang and Rust (2021) argue that AI increasingly complements or substitutes human intelligence across analytical, intuitive, and empathetic marketing tasks, altering the character of customer–firm interaction. This trajectory has moved through recognizable stages: rule-based automation for segmentation and email workflows (Haenlein & Kaplan, 2019); machine learning for prediction, recommendation, and churn management (Ma & Sun, 2020; Erevelles et al., 2016); and, most recently, Generative AI built on transformer-based Large Language Models (LLMs), which can produce text, images, code, and multimedia content and sustain natural-language interaction (Bommasani et al., 2021; Dwivedi et al., 2023).

Generative AI has measurably expanded what marketing technology can do — automated content production, conversational commerce, personalized advertising, and AI-assisted market research (Dwivedi et al., 2023; Mariani et al., 2022). Yet these systems remain fundamentally prompt-driven: a human specifies the objective, evaluates the output, and decides what happens next. Agentic Artificial Intelligence (Agentic AI) is presented in the emerging literature as the next stage in this evolution, integrating reasoning, planning, persistent memory, and tool use so that software agents can decompose a goal into subtasks, act through external systems, and adapt with comparatively little continuous human direction (Wang et al., 2024; Xi et al., 2023; Park et al., 2023). Industry and consultancy reports describe rapid experimentation with such agents across campaign orchestration, customer service, and sales workflows (Kim, 2025; Murugesan, 2025), but these accounts are not, by themselves, peer-reviewed evidence of marketing performance.

The interest in Agentic AI has therefore outpaced the scholarship needed to evaluate it. Reviews of AI in marketing have concentrated on predictive analytics and Generative AI (Dwivedi et al., 2023; Vrontis et al., 2022; Huang & Rust, 2021), and closely related work on AI agents in personal selling and customer experience is only beginning to accumulate (Ledro et al., 2022; sales-agent research reviewed in Section 4). A parallel and more skeptical literature has emerged around the governance risks of increasing machine autonomy, cautioning that fully autonomous systems raise unresolved questions of accountability, explainability, and control (Mitchell et al., 2025; Murugesan, 2025). The central problem addressed here is therefore not whether marketing organizations are experimenting with agentic systems — they plainly are — but how, and under what organizational and governance conditions, agentic AI affordances are converted into legitimate, repeatable marketing capability rather than into ungoverned automation or unsubstantiated efficiency claims.

1.1 Research gap and questions

The literature presents four connected gaps. First, much of the discourse treats AI autonomy itself as the explanatory variable, describing what agents can technically do while saying comparatively little about the organizational routines — data governance, role redesign, human oversight — needed to convert that capability into marketing value. Second, adoption, continued use, workflow change, and verified benefit are frequently collapsed into a single claim of “AI-driven performance,” obscuring which stage of implementation is actually being described. Third, the empirical base is thin and skewed: peer-reviewed, audited evidence of agentic-system performance in marketing is rare relative to the volume of conceptual papers, opinion pieces, and vendor or consultancy material in circulation. Fourth, responsible-AI governance for autonomous marketing agents — explainability, accountability, data protection across jurisdictions

such as the EU's GDPR and India's Digital Personal Data Protection regime — is typically discussed as an ethical add-on rather than as a condition for legitimate deployment.

The review addresses the following questions:

- RQ1: How has AI in marketing evolved from predictive and generative capability toward agentic autonomy in the scholarly literature?
- RQ2: Which organizational capabilities and governance conditions enable or constrain the conversion of Agentic AI into marketing value?
- RQ3: What can be concluded — and what cannot yet be concluded — about Agentic AI's marketing performance from the available evidence?
- RQ4: What theoretically grounded framework can guide implementation and future empirical testing?

1.2 Novelty and contribution

The contribution is more specific than pairing Generative AI with Agentic AI as sequential buzzwords. Theoretically, the review applies dynamic capabilities theory to explain how organizations sense which marketing decisions are suited to delegation, seize agentic affordances through bounded pilots, and reconfigure workflows, roles, and governance around them. Conceptually, it positions organizational capability — not agent autonomy alone — as the mechanism that converts AI affordances into marketing outcomes. Methodologically, it separates the small body of directly relevant, peer-reviewed evidence from the much larger volume of conceptual and industry commentary, and is explicit about that imbalance rather than citing vendor claims as though they were audited results. Practically, it proposes a staged capability framework in which data governance, human oversight, and accountability are designed in from the first pilot rather than added after deployment.

2. Review Method

2.1 Rationale for a critical narrative review

This article adopts a critical narrative review because its purpose is explanatory and theory-building: to connect research streams that use different vocabularies and methods (marketing, information systems, and autonomous-agent computer science), evaluate the limits of their claims, and develop an implementation framework. The design is informed by the quality domains in the Scale for the Assessment of Narrative Review Articles (SANRA) (Baethge et al., 2019).

A systematic literature review following PRISMA would be the appropriate design for an exhaustive, protocol-led answer to a narrowly bounded question with a stable, well-defined literature. Agentic AI in marketing does not yet meet that condition: terminology is unsettled, the peer-reviewed empirical base is small, and much of the visible discourse originates from preprints, conference proceedings, and industry publications rather than indexed journal databases that would support duplicate independent screening and a registered protocol. Claiming full PRISMA compliance without institutional access to bibliographic databases such as Scopus or Web of Science, and without a pre-registered protocol or independent dual screening, would overstate the reproducibility of the search. The critical narrative format is therefore the better fit: it allows the review to be transparent about what was actually done, to weigh conceptual and empirical sources differently, and to distinguish evidence from proposition without implying an exhaustiveness the search process cannot support.

2.2 Google Scholar search and source identification

Google Scholar was the principal platform used to identify literature, selected for its cross-disciplinary coverage of marketing, information systems, and computer science venues, including conference proceedings and preprints where a field is moving quickly. Searches were conducted during July–August 2026, with the final search performed on 6

August 2026. Its use is reported as a practical discovery strategy, not as evidence of exhaustive database coverage; publisher pages and DOI records were consulted only to verify bibliographic details after a source had been identified.

The search was iterative, combining exact phrases, synonyms, and contextual terms with Boolean operators. The principal search strings were:

- (“agentic AI” OR “AI agents” OR “autonomous AI”) AND (marketing OR advertising OR “customer experience”);
- (“generative AI” OR “large language model” OR LLM) AND (marketing OR advertising);
- (“artificial intelligence” OR AI) AND (marketing OR “customer relationship management”);
- (“AI agent” OR “autonomous agent”) AND (sales OR “customer service” OR “customer journey”);
- (“responsible AI” OR “explainable AI” OR accountability) AND (marketing OR “AI agents”); and
- (“dynamic capabilities” OR “AI capability”) AND (marketing OR firm performance), together with targeted searches connecting “multi-agent systems” and “human-AI collaboration” to marketing and service contexts.

The review began with all 35 reference entries in the source manuscript. Each was checked for a traceable title, complete authorship, outlet, year, and DOI or an equivalent stable identifier. All 35 were retained as bibliographically traceable and topically relevant; none were found to be fabricated, duplicated, or unrelated to the review questions. Google Scholar searching and citation chaining then identified 8 supplementary sources published through 2026 that were judged necessary to represent the specifically agentic (as opposed to generative) AI literature and the emerging governance debate around autonomous marketing systems, producing a cited corpus of 43 scholarly works. Backward searching examined the reference lists of retained sources; forward searching used Google Scholar’s “Cited by” function to locate later extensions and critiques.

2.3 Eligibility and structured review process

The review followed a structured narrative process of literature identification, relevance screening, critical appraisal, thematic coding, and narrative synthesis. It does not claim PRISMA compliance. Titles and available abstracts were first assessed against the review questions; potentially relevant works were then examined in full.

Eligible sources were English-language peer-reviewed journal articles; books and technical reports from established academic or standards publishers; and, given the pace of development in this specific sub-field, a small number of widely cited preprints (e.g., arXiv) and editorial or call-for-research pieces in peer-reviewed marketing journals where they represented the only available scholarly treatment of a specific claim. Each such preprint or editorial is explicitly identified as such at first mention rather than presented with the same evidentiary weight as a peer-reviewed empirical study. Blogs, vendor white papers, and unreviewed industry reports were not treated as scholarly evidence; where industry figures are mentioned in Section 4, they are attributed to their source and flagged as unaudited.

Table 1: Transparent review protocol

Element	Decision
Review type	Critical narrative review focused on theory development, conceptual integration, gap identification, and framework construction; no PRISMA claim.
Discovery platform	Google Scholar, supplemented by backward reference searching and Google Scholar forward-citation searching.

Element	Decision
Seed material	35 references in the uploaded manuscript; all 35 retained after traceability and relevance checks.
Search period	July–August 2026; final search completed on 6 August 2026.
Supplementary identification	Google Scholar searches and citation chaining; 8 sources added to the retained seed literature.
Time coverage	Primarily 2015–6 August 2026, with earlier seminal theory and foundational technical work retained.
Inclusion	English-language peer-reviewed articles; established academic books and technical reports; a small, explicitly flagged set of highly cited preprints and journal editorials where no peer-reviewed alternative existed.
Exclusion	Non-English publications; vendor white papers and unreviewed industry reports treated as evidence (rather than as flagged context); blogs; duplicates; untraceable citations; studies unrelated to marketing, AI capability, or agent governance.
Appraisal	Bibliographic traceability, fit between method and claim, distinction between conceptual and empirical contribution, and attention to limitations or ethics.
Extraction	Context, method, sample, theory, constructs, major findings, limitations, research gap, and managerial implications.
Synthesis	Thematic comparison organized by AI evolution, Agentic AI capability, marketing applications, governance, and evidence quality.

2.4 Critical appraisal, coding, and synthesis

No single appraisal instrument suits a corpus spanning conceptual articles, systematic reviews, technical surveys, preprints, and a handful of empirical studies. The review therefore used design-sensitive appraisal: empirical studies were assessed for sample transparency, context, measurement, and causal overreach; reviews and surveys were assessed for search transparency and the relationship between included studies and conclusions; conceptual and technical work was assessed for construct clarity and explanatory value. Preprints and editorial pieces were retained only where they represented a specific, otherwise-unavailable claim, and are flagged accordingly rather than silently blended with peer-reviewed findings.

Data were extracted into a comparison matrix and coded around five questions: what technology or capability was examined; what mechanism was proposed; what outcome, if any, was reported; what evidentiary status the source held (peer-reviewed empirical, peer-reviewed conceptual, preprint, or industry report); and what inference was justified. Codes were compared across five themes — AI evolution, Agentic AI capability, marketing applications, governance and accountability, and evidence quality — before being integrated narratively. The synthesis deliberately separates association from causation, vendor claims from audited findings, and reviewed evidence from the propositions developed later in the article.

3. Theoretical Foundation: Dynamic Capabilities

The resource-based view (RBV) treats valuable, difficult-to-imitate resources as a source of sustained advantage (Barney, 1991) and has been applied directly to AI adoption: Krakowski et al. (2023) show, using a chess-tournament setting, that AI adoption triggers interrelated substitution and complementation dynamics that can erode some human competitive capabilities while creating new ones. This is useful for explaining why analytical talent, creative expertise, and customer trust remain valuable marketing resources, but the RBV is less specific about how an organization renews those resources as agentic technology changes rapidly.

AI-capability research offers a complementary, more operational lens. Mikalef and Gupta (2021) conceptualize organizational AI capability as a bundle of tangible, human, and intangible resources that must be deliberately built rather than purchased, and Mishra et al. (2022) find that a firm's AI focus is associated with performance primarily where complementary organizational resources are present — an early empirical caution against assuming that AI adoption alone raises marketing performance. Neither of these perspectives, however, explains the dynamic, ongoing renewal process required when the technology itself — moving from predictive models to generative systems to autonomous agents — keeps changing the nature of the capability required.

Dynamic capabilities theory addresses this more directly. It defines higher-order capabilities through which organizations sense changes, seize opportunities, and reconfigure assets and routines (Teece, 2007). Applied to marketing, this logic implies that Agentic AI is not a one-time technology purchase but an iterative process in which an organization learns which marketing decisions are suitable for delegation to an autonomous agent, which data and integration infrastructure that delegation requires, where human judgment and brand stewardship remain indispensable, and how roles and controls should change as a result. In this review, sensing includes identifying which marketing tasks are repetitive, data-rich, and reversible enough to be candidates for agentic execution. Seizing includes building interoperable customer and campaign data, developing prompt- and agent-management expertise, and piloting bounded use cases with clear success and rollback criteria. Transforming includes redesigning campaign workflows, decision rights, creative review processes, and escalation paths so that autonomous execution is monitored rather than merely switched on.

The theory also guards against a technocentric reading of the evidence. Kunz et al. (2025) argue that human-machine collaboration in service and marketing contexts should be understood as an ongoing negotiation of roles rather than a fixed division of labor, and Murugesan (2025) similarly frames the rise of agentic systems as posing organizational and governance questions alongside technical ones. Dynamic capability, in this sense, is both technical and social: the same agentic technology can be seized productively by one organization and produce ungoverned automation in another, depending on the surrounding capability.

4. Critical Literature Synthesis

4.1 From predictive AI to generative AI in marketing

The trajectory from rule-based systems to machine learning to Generative AI is well established in the literature and is not seriously contested. Rule-based systems automated segmentation and email workflows but could not adapt beyond predefined logic (Haenlein & Kaplan, 2019). Machine learning extended AI into prediction — customer lifetime value, churn, demand forecasting, and recommendation — while strategic decisions remained with human managers (Ma & Sun, 2020; Erevelles et al., 2016; Zhang & Lu, 2021). Deep learning subsequently extended these capabilities to unstructured data such as text, images, and voice (LeCun et al., 2015).

Generative AI, built on transformer architectures and LLMs, is the most recent and best-evidenced stage: it supports automated content creation, conversational engagement, personalized advertising, and AI-assisted market research, and organizations report substantial reductions in content-development time (Dwivedi et al., 2023; Bubeck et al., 2023; Mariani et al., 2022). The literature is nonetheless consistent that Generative AI remains fundamentally prompt-

driven — a human specifies the objective, evaluates output, and decides the next action — and that its documented limitations (hallucination, bias, homogenization of creative output) are limitations of a system that assists rather than acts (Dwivedi et al., 2023).

Synthesis. This stage of the evolution is comparatively well supported by peer-reviewed evidence and is not the primary source of uncertainty in the current debate. The more consequential and less settled question, addressed next, is what changes — technologically and organizationally — once a system moves from generating content on request to acting on an objective with limited supervision.

4.2 Agentic AI: concept, architecture, and the limits of current evidence

Agentic AI is defined in the technical literature by the combination of reasoning, planning, persistent memory, tool use, and autonomous execution within a single architecture (Wang et al., 2024; Xi et al., 2023). Park et al. (2023) demonstrate, in a simulated social environment rather than a commercial marketing setting, that generative agents can plan and coordinate believable multi-step behavior over time; Schick et al. (2023) show that language models can learn to invoke external tools and APIs, which is the technical basis for agents that can query a CRM system, place an order, or launch a campaign rather than only describe one. These are foundational computer-science results, not marketing studies, and this distinction matters: much of the marketing literature's enthusiasm for Agentic AI rests on extrapolating from these general-purpose capabilities rather than on marketing-specific evidence.

Kim (2025), writing an editorial call for research in the *Journal of Interactive Advertising*, argues that agentic systems are already reshaping the advertising process by shifting interaction from search to dialogue and from advertiser-placed messages toward answer-engine-mediated exchanges, and calls for theoretical adaptation because classic persuasion models assume a human-authored, human-placed advertisement. This is a credible and influential early academic statement of the problem, but it is explicitly a call for research rather than a report of findings, and it should be read as agenda-setting rather than as evidence of measured effects.

On the governance side, Mitchell et al. (2025) argue directly against the pursuit of fully autonomous AI agents, contending that meaningful human oversight and correctability should remain a design requirement rather than an incidental feature; Murugesan (2025) reaches a related conclusion, cataloguing the accountability, safety, and control questions raised by increasing agent autonomy. These positions sit in direct tension with vendor and consultancy narratives that treat autonomy itself as the desirable end-state of marketing automation, and the review treats that tension as a genuine, unresolved disagreement in the literature rather than resolving it in either direction.

Synthesis. The technical foundations of Agentic AI are reasonably well documented, but marketing-specific evidence of what autonomous agents do differently, and better, than generative or predictive systems is early and largely conceptual. The field currently has clearer evidence of what agentic systems can technically do than of what they should be trusted to do unsupervised in a marketing context.

4.3 Applications across the customer journey

Within this early evidence base, the literature converges on a recognizable set of candidate applications rather than a single dominant use case. In customer relationship management, autonomous agents are described as monitoring behaviour, anticipating needs, and initiating retention actions with reduced continuous supervision (Ma & Sun, 2020; Wirtz et al., 2018). Ledro et al. (2022), reviewing AI in customer relationship management more broadly, find that the literature consistently reports operational gains in service efficiency and personalization but rarely isolates the specific contribution of autonomy versus simpler automation — a caution that applies directly to more recent agentic claims. In digital advertising and campaign management, agents are proposed to allocate budget, run experiments, and adjust targeting in response to performance signals (Wang et al., 2024; Kim, 2025).

The clearest empirical signal currently available comes from an adjacent function, personal selling, where autonomous agents are already commercially deployed for prospecting and lead qualification. Recent sales research finds that

human salespeople remain more effective than AI agents at relationship-building tasks in exploration and expansion phases, with psychological distance as an explanatory mechanism, while buyers with higher technological proficiency tend to prefer interacting with agents — evidence that the appropriate allocation of tasks between humans and agents is conditional on task type and customer segment rather than uniformly favouring autonomy. Because this evidence originates in sales rather than marketing proper, it should be read as a plausible analogue, not direct proof of marketing outcomes.

Mishra et al. (2022) provide one of the few firm-level empirical studies connecting AI to marketing-relevant performance, finding that a firm's AI focus is associated with financial and market performance, though the relationship is conditioned by complementary capabilities and is not specific to agentic or autonomous systems. Dynamic-pricing, market-intelligence, and brand-monitoring applications appear frequently in conceptual and industry discussion (Wang et al., 2024) but rest on a comparatively thin peer-reviewed base at the time of this review.

Synthesis. Plausible marketing applications for Agentic AI span the customer journey, but the strength of supporting evidence varies sharply by function: adjacent-domain evidence in personal selling is the most directly comparable, firm-level performance evidence is sparse and not agent-specific, and applications such as autonomous dynamic pricing and campaign orchestration are argued conceptually more often than they are demonstrated empirically.

4.4 Governance, accountability, and responsible AI in autonomous marketing

Responsible-AI scholarship converges on fairness, transparency, accountability, privacy, and human oversight as necessary conditions for legitimate deployment, but disagrees on how to operationalize them once a system can act rather than only recommend (Dignum, 2019; Floridi & Cows, 2019). For Agentic AI specifically, the core governance problem is that autonomy compresses the interval between a system's output and its real-world consequence: a Generative AI hallucination produces a bad draft that a human can catch before publication, whereas an agentic system with tool access can place a bid, send a message, or adjust a price before a human reviews it (Dwivedi et al., 2023; Mitchell et al., 2025).

This has direct implications for explainability and accountability in marketing decisions. When an autonomous agent's pricing or targeting decision produces an undesirable or discriminatory outcome, responsibility may lie with the underlying model, the data it was trained or grounded on, the tool permissions it was granted, or the absence of a human checkpoint — and the literature does not yet offer a settled way to apportion that responsibility once execution, not just recommendation, has been delegated. Kunz et al. (2025) argue that human oversight should be understood as an evolving collaborative role — human as strategist and governor rather than operator — rather than a static checkpoint that autonomy is expected to eventually eliminate. Mitchell et al. (2025) go further, arguing that some categories of fully autonomous deployment should not be pursued at all until correctability and oversight mechanisms are demonstrably reliable.

Cross-jurisdictional data-protection obligations compound this problem for marketing specifically, since agentic systems typically require continuous access to customer data to personalize and act. Organizations operating in India must design around the Digital Personal Data Protection Act 2023 and its 2025 implementing rules, while organizations with European customers remain subject to the General Data Protection Regulation; neither regime was written with autonomous, tool-using marketing agents specifically in mind, which leaves compliance interpretation to individual organizations in the near term.

Synthesis. Governance is not a compliance layer to be added after an agentic system is built; the literature's most autonomy-skeptical voices treat human oversight, correctability, and jurisdiction-appropriate data governance as design requirements for legitimate deployment. Any marketing capability framework that treats autonomy as an unambiguous improvement over supervised generation is out of step with this part of the literature.

4.5 Evidence quality and inferential boundaries

The evidence base for Agentic AI in marketing is wide in circulation but shallow in peer-reviewed depth. A large share of the material available through general web and even scholarly search consists of consultancy reports, vendor blog posts, and trade-press articles reporting adoption percentages, productivity gains, and return-on-investment estimates. These figures are not independently audited, rarely disclose sampling or measurement methods, and are commercially interested in favourable framing; this review does not treat them as evidence of marketing performance, and cites them, where mentioned at all, strictly as indicators of industry discourse rather than as findings.

Within the peer-reviewed literature itself, most Agentic-AI-specific work remains conceptual, editorial, or technical-computer-science in orientation rather than empirical marketing research (Wang et al., 2024; Xi et al., 2023; Kim, 2025). The empirical studies that do exist are concentrated in adjacent domains — personal selling, customer service — or examine AI focus and firm performance without isolating agentic autonomy as the mechanism (Mishra et al., 2022). Longitudinal studies, comparison groups, natural experiments, and algorithm audits specific to marketing agents are, at the time of this review, essentially absent from the peer-reviewed record. This imbalance is itself a finding: the pace of technological and commercial deployment is running well ahead of the pace of independent evaluation, and any framework built from this literature must be presented as a synthesis of plausible mechanisms and open questions rather than as a summary of demonstrated effects.

Table 2: Critical literature review matrix

Author/year	Evidentiary status	Focus	Major finding	Limitation / gap
Huang & Rust (2021)	Peer-reviewed conceptual framework	AI's role across analytical, intuitive, and empathetic marketing tasks	AI increasingly complements or replaces human intelligence in stages, reshaping customer-firm interaction.	Conceptual; predates the specifically agentic literature.
Davenport et al. (2020)	Peer-reviewed conceptual review	AI's effect on the future of marketing	AI shifts marketing from support to strategic capability but requires organizational change to realize value.	Conceptual synthesis; no agent-specific evidence.
Wang et al. (2024)	Peer-reviewed technical survey	LLM-based autonomous agents	Documents architectures for planning, memory, and tool use that underpin agentic systems.	Computer-science survey; not marketing-specific.
Park et al. (2023)	Peer-reviewed empirical (simulation)	Generative agents in a simulated social environment	Agents can plan and sustain believable multi-step behaviour over simulated time.	Simulated, non-commercial setting; no marketing outcome measured.
Schick et al.	Peer-reviewed	Tool-use learning in	Models can learn to	Technical

Author/year	Evidentiary status	Focus	Major finding	Limitation / gap
(2023)	empirical (technical)	language models	invoke external tools/APIs, enabling agent action beyond text generation.	benchmark study; not evaluated in a marketing context.
Krakowski et al. (2023)	Peer-reviewed empirical	RBV and AI adoption using chess as a controlled setting	AI adoption triggers substitution and complementation dynamics that reshape competitive capability.	Controlled non-marketing setting; transfer to marketing is inferential.
Mishra et al. (2022)	Peer-reviewed empirical	Firm AI focus and performance	AI focus is associated with firm performance, conditioned by complementary organizational resources.	Firm-level AI focus, not agentic autonomy specifically.
Ledro et al. (2022)	Peer-reviewed literature review	AI in customer relationship management	AI improves CRM efficiency and personalization, but the literature rarely isolates autonomy's specific contribution.	Predates the current agentic-AI wave.
Kim (2025)	Peer-reviewed editorial / call for research	Agentic AI and advertising theory	Agentic, dialogue-based AI requires rethinking classic persuasion and media-placement theory.	Explicitly agenda-setting, not an empirical finding.
Mitchell et al. (2025)	Preprint (flagged)	Governance limits of autonomous agents	Argues fully autonomous agents should not be developed without reliable correctability and oversight.	Position paper; not marketing-specific, but directly relevant to governance.
Murugesan (2025)	Peer-reviewed commentary	Rise of agentic AI: implications and concerns	Identifies accountability, safety, and control as unresolved organizational questions accompanying agent autonomy.	Cross-domain commentary rather than marketing field study.

5. Review Findings

5.1 Finding 1: A conversion chain, not a direct autonomy effect

Across the reviewed literature, the most defensible sequence is: Generative AI affordances expand what content and interaction an organization can produce; Agentic AI affordances expand what an organization can execute without continuous supervision; organizational capability — data, skills, and workflow — converts that execution capacity into marketing action; and governance determines whether the result is accountable and sustainable. Treating agent autonomy as a direct predictor of marketing performance, without specifying this conversion chain, is the same conceptual error identified in adjacent literatures that study AI and human resource management: technology adoption is not itself the mechanism of value.

5.2 Finding 2: Organizational readiness is multidimensional and function-specific

Readiness for agentic delegation includes interoperable customer and campaign data, staff able to design, monitor, and correct agent behaviour, defined escalation and rollback procedures, and leadership willing to fund evaluation rather than only deployment. A firm may be ready to let an agent draft and A/B test ad copy while being nowhere near ready to let one set live prices or approve creative for a sensitive campaign. Readiness should therefore be assessed at the use-case level rather than as a single organizational attribute.

5.3 Finding 3: Benefits are plausible but the evidence base is early and imbalanced

The literature consistently proposes efficiency, personalization, and responsiveness benefits from agentic marketing systems, but the supporting evidence is dominated by conceptual papers, technical computer-science studies conducted outside marketing settings, and unaudited industry figures. Peer-reviewed, marketing-specific, outcome-audited studies of agentic systems are rare. The review therefore supports cautiously worded expectations about Agentic AI's marketing value, not claims that autonomous agents have been shown to significantly improve marketing performance.

5.4 Finding 4: Autonomy sharpens, rather than removes, the case for governance

The most autonomy-skeptical sources in this review (Mitchell et al., 2025; Murugesan, 2025) do not argue against Agentic AI in general; they argue that increasing autonomy raises the stakes of inadequate oversight because the interval between a system's decision and its real-world consequence shrinks. This directly qualifies more optimistic industry and conceptual accounts that treat autonomy as an unambiguous efficiency gain.

5.5 Finding 5: Cross-jurisdictional and adjacent-domain evidence is being used to fill a marketing-specific gap

In the absence of marketing-specific longitudinal or audited evidence, the literature leans heavily on adjacent-domain findings (personal selling, customer service) and general firm-level AI-performance studies. These sources are informative about plausible mechanisms but cannot establish effect size, feasibility, or legitimacy specifically for marketing decisions, nor can they resolve how agentic systems should be governed differently under, for example, India's data-protection regime versus the EU's. This is the review's clearest evidence gap.

6. Proposed Conceptual Framework: Evolution Toward Autonomous Marketing Ecosystems

Building on the critical synthesis above, the framework treats AI evolution in marketing as a progressive increase in intelligence, autonomy, and decision-making scope — while explicitly reintroducing the organizational-capability and governance conditions that Section 4 shows are necessary for that progression to translate into legitimate marketing value. The framework distinguishes four evolutionary stages: Predictive AI, Generative AI, Agentic AI, and Autonomous Marketing Ecosystems.

Table 3: Evolution of AI toward autonomous marketing ecosystems

Evolution stage	Core capability	Marketing role	Major applications	Governance requirement
Predictive AI	Data analysis, forecasting, pattern recognition	Decision-support system	Customer segmentation, churn prediction, recommendation systems	Data quality and model validation
Generative AI	Content generation and conversational intelligence	Creative and analytical assistant	Content marketing, chatbots, personalized communication, marketing research	Human review before publication; disclosure of AI-generated content
Agentic AI	Autonomous reasoning, planning, memory, and execution	Bounded autonomous collaborator	Campaign optimization, lead qualification, market intelligence	Defined tool permissions, rollback procedures, human escalation points
Autonomous marketing ecosystem	Multi-agent collaboration and self-optimization	Independently operating marketing function	AI-managed customer relationships, autonomous negotiation and transactions	Continuous audit, accountability assignment, jurisdiction-specific compliance

The framework departs from a purely technological reading by treating the distinction between Generative AI and Agentic AI as a distinction in autonomy and accountability, not only intelligence. Generative AI augments human creativity and productivity under continuous human evaluation; Agentic AI shifts execution toward the system itself, which is precisely why Section 4.4 identifies governance as a design requirement rather than a later safeguard.

Table 4: Comparison between Generative AI and Agentic AI in marketing

Dimension	Generative AI	Agentic AI
Primary objective	Content generation and knowledge creation	Autonomous pursuit of a defined marketing goal
Decision-making	Reactive response to prompts	Proactive, multi-step decision execution
Human role	Prompt creation and output evaluation	Strategic supervision, permissioning, and governance
Planning ability	Limited to the current prompt/response cycle	Multi-step planning and sequencing
Memory	Short-term contextual memory	Persistent organizational and customer memory

Dimension	Generative AI	Agentic AI
Learning	Limited within-session adaptation	Continuous adaptation through feedback and outcomes
Tool integration	Limited external system interaction	Direct integration with CRM, ERP, APIs, and databases
Accountability locus	Clear - rests with the human evaluating output	Distributed across model, data, permissions, and oversight design

6.1 The AI–Capability–Marketing relationship

The framework clarifies the distinct roles of AI technology, organizational capability, and marketing practice, extending the logic used in Section 3. Agentic AI techniques operate on customer and market data to plan, act, and adapt. Organizational capability — data infrastructure, agent-management skill, defined permissions, and escalation routines — converts that technical capacity into governed marketing action. Marketing practice is the organizational site where a campaign, price, or customer interaction is actually executed and experienced by a customer. Outcome and incident evidence then feeds back into capability and governance, closing a learning loop rather than a one-time deployment. Collapsing these layers into a single “AI adoption” variable, as much of the industry discourse does, obscures accountability: when an agentic decision produces an undesirable outcome, the cause may lie in the model, the data, the permission scope granted to the agent, or the absence of a human checkpoint at the relevant stage.

7. Proposed Research Model and Propositions

The propositions below are offered for future empirical testing; they are not presented as tested results. AI capability is defined as an organization’s ability to select, integrate, and operate reliable Generative and Agentic AI applications; marketing capability is the ability to convert that technical capacity into governed, evaluable marketing practice; and outcomes are multi-level rather than a single performance score.

P1. AI capability will be positively associated with realized marketing value only where organizational capability — data infrastructure, agent-management skill, and workflow redesign — is also present; AI capability alone is expected to show a weak or inconsistent direct association with marketing outcomes.

P2. The relationship between Agentic AI adoption and marketing performance will be mediated by the degree of workflow redesign (decision rights, escalation routes, and review procedures) rather than by autonomy adoption itself.

P3. Task reversibility and data quality will moderate the relationship between agent autonomy and marketing performance: low-risk, reversible, data-rich tasks (e.g., ad-copy testing) will show a more favourable autonomy–performance relationship than high-stakes, low-reversibility tasks (e.g., live pricing or brand-sensitive messaging).

P4. Responsible-governance mechanisms (defined tool permissions, human escalation points, and audit routines) will strengthen positive outcomes and reduce adverse outcomes from agentic deployment by improving error detection and accountable correction.

P5. Organizational learning-orientation and leadership support will strengthen the relationship between AI capability and marketing-practice redesign, consistent with dynamic-capabilities theory’s emphasis on reconfiguration as an ongoing process rather than a one-time adoption event.

P6. Customer technological proficiency and task type will moderate customer acceptance of agent-mediated interactions, consistent with adjacent sales-agent evidence that more technologically proficient customers respond more favourably to agent interaction than others.

P7. Outcome monitoring and incident feedback will support continuing reconfiguration of agentic marketing capability, producing a learning loop rather than a one-time performance effect — and organizations that treat autonomous deployment as reversible will detect and correct errors faster than those that do not.

8. Discussion

8.1 Theoretical implications

The review extends AI-marketing theory by locating value creation in a capability-conversion process rather than in agent autonomy itself. The resource-based view and AI-capability research (Barney, 1991; Mikalef & Gupta, 2021; Mishra et al., 2022) explain why complementary organizational resources matter; dynamic capabilities theory explains why some organizations can repeatedly reconfigure those resources as agentic technology evolves while others remain at the pilot stage. AI capability is a lower-order enabling capability; organizational (data, skill, workflow) capability is the bridging mechanism; governance and human oversight determine whether the resulting routines are legitimate and sustainable rather than merely fast.

The framework also qualifies universalistic claims about autonomy. A marketing practice is not automatically improved by being delegated to an autonomous agent; its value depends on task reversibility, data validity, customer segment, and the consequences of error — a position directly supported by the tension between optimistic industry accounts and the autonomy-skeptical positions of Mitchell et al. (2025) and Murugesan (2025).

8.2 Methodological implications

Future studies should avoid inferring marketing performance from general AI-adoption measures or from vendor-reported productivity figures. Multi-source designs should distinguish AI capability, actual agent use and permission scope, workflow change, and independently measured marketing outcomes. Where feasible, matched comparisons, stepped-wedge pilots, and algorithm audits specific to marketing agents would substantially strengthen causal inference relative to the current cross-sectional, conceptual, and vendor-reported base identified in Section 4.5. Measurement should also capture potential harms — brand-reputation incidents, discriminatory targeting, and customer trust erosion — alongside efficiency benefits.

8.3 Practical implications

The implications differ by stakeholder role and authority over the system. Table 5 converts the review’s findings into role-specific priorities.

Table 5: Stakeholder-specific recommendations

Stakeholder	Priority recommendations	Minimum safeguard before scaling
Marketing leadership (CMO and equivalent)	Govern a defined portfolio of marketing decisions suitable for agentic delegation rather than adopting autonomy as a blanket strategy; fund evaluation and staff learning alongside deployment.	Documented use-case scope, accountable owner, and authority to pause or retire an underperforming agent.
Marketing and analytics teams	Build reliable, integrated customer and campaign data; separate descriptive, predictive, and agentic-execution claims; design workflows with explicit human checkpoints.	Data lineage, defined agent permissions, documented escalation and override procedures.

Stakeholder	Priority recommendations	Minimum safeguard before scaling
Customers and employee representatives	Retain visibility into when an interaction is agent-mediated; report concerns about targeting, pricing fairness, or loss of human contact through accessible channels.	Disclosure of AI-mediated interaction where material; accessible correction and complaint routes.
Policymakers and regulators	Clarify how existing data-protection and consumer-protection regimes (e.g., India's DPDP Act, the EU's GDPR) apply to autonomous, tool-using marketing agents specifically.	Guidance distinguishing low-risk decision support from high-stakes autonomous execution.
Technology vendors	Disclose intended and excluded uses, validation evidence, and the limits of agent autonomy granted by default configurations.	Change notification, audit logging, and a documented human-override mechanism.

9. Limitations

This article is a critical narrative review and is not exhaustive. Google Scholar offers broad interdisciplinary discovery but has opaque and changing coverage, ranking, and citation-linking practices, so searches performed at another date or by another reviewer may return a different set of records. The review did not use subscription-database exports, independent dual screening, a registered protocol, or meta-analysis; selection and interpretation retain an element of reviewer judgment despite the stated criteria and source accounting.

Because Agentic AI in marketing is an emerging field, the corpus necessarily includes a small number of preprints and an editorial call-for-research piece alongside peer-reviewed empirical and conceptual work; each is explicitly flagged where cited, but their inclusion means the review's synthesis is, in places, built on claims that have not yet passed peer review in their current form. Restriction to English-language sources creates language bias and may omit relevant regional scholarship, including Indian marketing-technology research not indexed in English-language venues. Backward and forward citation searching can introduce citation bias by favouring established, highly cited work over newer or less-connected studies.

The empirical evidence directly available on Agentic AI in marketing is sparse; much of the applied reasoning in this review draws on adjacent domains (personal selling, customer service, general AI-firm-performance research) and on technical computer-science studies conducted outside commercial marketing settings. These sources can identify plausible mechanisms but cannot establish effect size, feasibility, or transferability to specific marketing functions or to the Indian regulatory and market context. The pace of both AI development and data-protection implementation means governance requirements and technical practice will continue to evolve after this review is completed.

These limitations are partly addressed through explicit search strings, eligibility criteria, source accounting, evidentiary-status labelling of each source, and restrained causal language. They cannot be eliminated by reporting alone. The proposed framework should be treated as a theoretically grounded synthesis and research agenda requiring empirical validation, not as a summary of demonstrated marketing performance effects.

10. Future Research Agenda

Future research should prioritize designs that can distinguish adoption from capability, vendor claims from audited outcomes, and association from causation. Table 6 sets out a structured agenda.

Table 6: Structured future research agenda

Research gap	Suggested future study	Suggested methodology	Expected contribution
Unknown prevalence and maturity of agentic deployment	Establish a baseline of formal, pilot, and informal Agentic AI use across marketing functions and firm sizes.	Cross-sectional survey combined with a system/tool inventory.	Distinguish adoption claims from operational capability and identify patterns in readiness and use cases.
Unobserved capability-conversion process	Follow bounded agentic pilots from problem definition through deployment, monitoring, and renewal or retirement.	Longitudinal case studies using interviews, system logs, and process tracing.	Test the sensing–seizing–transforming mechanism and explain why similar technology produces different outcomes.
Weak causal evidence of marketing benefit	Evaluate a clearly defined, reversible agentic intervention (e.g., agent-managed ad-copy testing) against an appropriate human baseline.	Matched comparison or stepped-wedge field trial where feasible and ethical.	Estimate whether observed changes follow the intervention rather than prior trends or selection.
Limited customer-centred evidence	Examine customer trust, disclosure preferences, and satisfaction when interactions are agent-mediated versus human-mediated.	Experimental or quasi-experimental consumer study across technological-proficiency segments.	Extend evaluation beyond firm-side efficiency to customer experience and trust.
Unverified fairness and accountability in autonomous execution	Audit agentic pricing, targeting, or recommendation systems for error distribution, disparate impact, and override patterns.	Socio-technical algorithm audit combining quantitative error analysis with workflow observation.	Operationalize responsible-AI governance for marketing agents and show whether safeguards function in practice.
Cross-jurisdictional governance gap	Compare how data-protection and consumer-protection regimes (India's DPDP Act, the EU's GDPR, and others) are being interpreted for autonomous marketing agents.	Comparative policy and documentation analysis supplemented by practitioner interviews.	Clarify the distribution of compliance responsibility across markets and inform proportionate regulatory guidance.

Research gap	Suggested future study	Suggested methodology	Expected contribution
Untested proposed model	Evaluate the relationships in Section 7 without relying on a single respondent, vendor-reported metric, or measurement occasion.	Multi-source, time-separated design using marketing, IT, governance, and customer data.	Test mediation and moderation while separating AI capability, workflow change, and independently verified outcomes.

11. Conclusion

Agentic AI can plausibly extend marketing capability beyond what Generative AI supports, but the current literature does not show that autonomy alone improves marketing outcomes. The most defensible reading of the evidence is conditional: Generative AI expands what an organization can create, Agentic AI expands what it can execute without continuous supervision, organizational capability converts that execution capacity into governed marketing action, and responsible-AI governance determines whether the result is legitimate and sustainable. Marketing-specific, peer-reviewed, and audited evidence remains too limited for claims of established performance effects, and much of the confidence currently circulating in industry discourse rests on unaudited vendor and consultancy figures rather than independently verified findings.

The practical implication is a shift from autonomy adoption to capability building. Organizations should begin with bounded, reversible marketing decisions, invest in reliable data and agent-management skill, design human checkpoints and escalation routes before scaling, and evaluate benefit and harm before extending an agent's scope. This sequence aligns agentic marketing experimentation with the accountability that customers, employees, and regulators reasonably expect, while creating an empirically testable agenda for the research that this fast-moving field still needs.

Declarations

Data availability. No new empirical dataset was created. The review is based on the sources listed in the reference section.

Ethics approval. Not applicable; the article does not report research involving human participants or identifiable personal data.

Funding. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of interest. The author declares no conflict of interest.

References

1. Bagozzi, R. P., & Dholakia, U. M. (2002). Intentional social action in virtual communities. *Journal of Interactive Marketing*, 16(2), 2–21. <https://doi.org/10.1002/dir.10006>
2. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., et al. (2021). On the opportunities and risks of foundation models. *arXiv*. <https://doi.org/10.48550/arXiv.2108.07258>
3. Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrmann, S., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., et al. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv*. <https://doi.org/10.48550/arXiv.2303.12712>

4. Davenport, T. H., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48, 24–42. <https://doi.org/10.1007/s11747-019-00696-0>
5. Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.
6. Dignum, V. (2019). *Responsible artificial intelligence: How to develop and use AI in a responsible way*. Springer.
7. Duan, Y., Edwards, J. S., & Dwivedi, Y. K. (2019). Artificial intelligence for decision making in the era of big data—Evolution, challenges and research agenda. *International Journal of Information Management*, 48, 63–71. <https://doi.org/10.1016/j.ijinfomgt.2019.01.021>
8. Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P. V., Janssen, M., Jones, P., Kar, A. K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., ... Williams, M. D. (2021). Artificial intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57, Article 101994. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>
9. Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., ... Wright, R. (2023). Opinion paper: “So what if ChatGPT wrote it?” Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, Article 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
10. Erevelles, S., Fukawa, N., & Swayne, L. (2016). Big data consumer analytics and the transformation of marketing. *Journal of Business Research*, 69(2), 897–904. <https://doi.org/10.1016/j.jbusres.2015.07.001>
11. Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.8cd550d1>
12. Haenlein, M., & Kaplan, A. (2019). A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. *California Management Review*, 61(4), 5–14. <https://doi.org/10.1177/0008125619864925>
13. Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*, 49, 30–50. <https://doi.org/10.1007/s11747-020-00749-9>
14. Huang, M.-H., & Rust, R. T. (2022). A framework for collaborative artificial intelligence in marketing. *Journal of Retailing*, 98(2), 209–223. <https://doi.org/10.1016/j.jretai.2021.10.001>
15. Kaplan, A., & Haenlein, M. (2019). Siri, Siri in my hand: Who’s the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15–25. <https://doi.org/10.1016/j.bushor.2018.08.004>
16. Kim, J. (2025). Advertising in the age of agentic AI: Call for research. *Journal of Interactive Advertising*, 25(3), 215–221. <https://doi.org/10.1080/15252019.2025.2557107>
17. Kitchenham, B., & Charters, S. (2007). *Guidelines for performing systematic literature reviews in software engineering* (Technical Report EBSE-2007-01). Keele University.
18. Krakowski, S., Luger, J., & Raisch, S. (2023). Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal*, 44(6), 1425–1452. <https://doi.org/10.1002/smj.3387>

19. Kshetri, N. (2024). Generative artificial intelligence in business and society: Opportunities and challenges. *IEEE Computer*, 57(1), 102–106.
20. Kunz, W. H., Sajtos, L., & Flavián, C. (2025). Beyond replacement: Human-machine collaboration in the age of AI. *Journal of Service Management*, 36, 477–494. <https://doi.org/10.1108/JOSM-04-2025-0194>
21. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444. <https://doi.org/10.1038/nature14539>
22. Ledro, C., Nosella, A., & Vinelli, A. (2022). Artificial intelligence in customer relationship management: Literature review and future research directions. *Journal of Business & Industrial Marketing*, 37(13), 48–63. <https://doi.org/10.1108/JBIM-07-2021-0332>
23. Ma, L., & Sun, B. (2020). Machine learning and AI in marketing—Connecting computing power to human insight. *International Journal of Research in Marketing*, 37(3), 481–504. <https://doi.org/10.1016/j.ijresmar.2020.04.005>
24. Mariani, M., Perez-Vega, R., & Wirtz, J. (2022). AI in marketing, consumer research and psychology. *Psychology & Marketing*, 39(4), 755–776. <https://doi.org/10.1002/mar.21653>
25. Mikalef, P., & Gupta, M. (2021). Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study. *Information & Management*, 58(3), Article 103434. <https://doi.org/10.1016/j.im.2020.103434>
26. Mishra, S., Ewing, M. T., & Cooper, H. B. (2022). Artificial intelligence focus and firm performance. *Journal of the Academy of Marketing Science*, 50, 1176–1197. <https://doi.org/10.1007/s11747-022-00876-5>
27. Mitchell, M., Ghosh, A., Luccioni, A. S., et al. (2025). Fully autonomous AI agents should not be developed. *arXiv*. <https://doi.org/10.48550/arXiv.2502.02649>
28. Murugesan, S. (2025). The rise of agentic AI: Implications, concerns, and the path forward. *IEEE Intelligent Systems*, 40(2), 8–14. <https://doi.org/10.1109/MIS.2025.3544940>
29. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>
30. Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. *Proceedings of the ACM Symposium on User Interface Software and Technology*, 1–22. <https://doi.org/10.1145/3586183.3606763>
31. Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
32. Rust, R. T. (2020). The future of marketing. *International Journal of Research in Marketing*, 37(1), 15–26. <https://doi.org/10.1016/j.ijresmar.2019.08.002>
33. Rust, R. T., & Huang, M.-H. (2021). The feeling economy: Managing in the next generation of artificial intelligence (AI). *California Management Review*, 63(4), 43–65. <https://doi.org/10.1177/00081256211017887>
34. Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36, 68539–68551.
35. Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104, 333–339. <https://doi.org/10.1016/j.jbusres.2019.07.039>

36. Syam, N., & Sharma, A. (2018). Waiting for a sales renaissance in the fourth industrial revolution: Machine learning and artificial intelligence in sales research and practice. *Industrial Marketing Management*, 69, 135–146. <https://doi.org/10.1016/j.indmarman.2017.12.019>
37. Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), 1319–1350. <https://doi.org/10.1002/smj.640>
38. Thomas, J., & Harden, A. (2008). Methods for the thematic synthesis of qualitative research in systematic reviews. *BMC Medical Research Methodology*, 8, Article 45. <https://doi.org/10.1186/1471-2288-8-45>
39. Tranfield, D., Denyer, D., & Smart, P. (2003). Towards a methodology for developing evidence-informed management knowledge by means of systematic review. *British Journal of Management*, 14(3), 207–222. <https://doi.org/10.1111/1467-8551.00375>
40. Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., & others. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18, Article 186345. <https://doi.org/10.1007/s11704-024-40231-9>
41. Wirtz, J., Patterson, P. G., Kunz, W. H., Gruber, T., Lu, V. N., Paluch, S., & Martins, A. (2018). Brave new world: Service robots in the frontline. *Journal of Service Management*, 29(5), 907–931. <https://doi.org/10.1108/JOSM-04-2018-0119>
42. Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, S., Zhang, M., Wang, J., Jin, B., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., & others. (2023). The rise and potential of large language model based agents: A survey. *arXiv*. <https://doi.org/10.48550/arXiv.2309.07864>
43. Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23, Article 100224. <https://doi.org/10.1016/j.jii.2021.100224>