

AI Adoption: How Small and Medium-Sized Enterprises (SMEs) Are Leveraging Generative AI for Business Decision-Making and Cost Reduction

Milind Gajghate¹, Bhairavee Bhagadkar²

^{1,2}Department of Management Studies
Guru Nanak Institute of Engineering & Technology, Nagpur

Peer Review Information	Abstract
Type: Article Received: 11 July 2026 Revised: 13 Aug2026 Accepted: 07 Sept 2026 Published: 06 Oct 2026	The democratization of artificial intelligence through generative AI (GenAI) has altered the technological options available to small and medium-sized enterprises (SMEs). This study examines how SMEs use GenAI to support business decision-making and reduce operating costs. Using the mixed-methods multi-case basis stated in the source document, covering 45 SMEs across manufacturing, professional services, and retail, the paper synthesizes functional adoption pathways in market intelligence, financial scenario analysis, customer support, marketing, administration, and software development. The reported evidence indicates that GenAI can scale cognitive capabilities while automating repetitive administrative, creative, and technical work. The paper also retains a cost-optimization model that relates initial AI integration expenditure and recurring infrastructure cost to reductions in marginal transaction cost. Across the reported three-year projections, aggregate functional spending declines from a baseline of USD 410,000 to USD 170,000 in the mature state, corresponding to the reported net saving of 58.5%. The study further identifies data confidentiality, hallucination, algorithmic bias, and vendor dependency as implementation risks and proposes a phased governance and adoption roadmap for SMEs. Keywords: Generative AI; Small and Medium-Sized Enterprises; Business Decision-Making; Cost Reduction; Digital Transformation; AI Governance

How to Cite This Article

Gajghate, M., & Bhagadkar, B. (2026). AI adoption: How small and medium-sized enterprises (SMEs) are leveraging generative AI for business decision-making and cost reduction. *International Journal of Research and Development: A Management Review*, 15(2), 588–592.

Introduction

Small and medium-sized enterprises form a major part of global business activity, yet their adoption of advanced digital technologies has historically been constrained by capital, infrastructure, and specialist skill requirements. The emergence of accessible large language models, multimodal systems, and application programming interfaces changes this cost structure by allowing firms to deploy natural-language-driven cognitive tools without building large proprietary machine-learning infrastructures. In the source study, this shift is framed as a movement from capital-intensive engineering toward prompt-centric and agile customization.

The managerial opportunity extends beyond content generation. GenAI can translate natural-language questions into analytical workflows, summarize complex operational information, generate scenarios, automate service interactions, and assist software development. For resource-constrained SMEs, these capabilities may improve the speed of decision-making while reducing the human effort attached to high-frequency tasks. At the same time, adoption introduces risks involving confidential data, inaccurate model outputs, bias, compliance, and dependence on external AI vendors.

Problem Statement and Research Gap

Although generative AI is being adopted rapidly, the source document identifies a gap in formal frameworks focused specifically on SME operating environments. Smaller firms have tighter working-capital cycles, lower tolerance for implementation failure, and shorter strategic horizons than large corporations. The practical research problem is therefore how SMEs can move from ad-hoc GenAI use toward structured decision support and cost reduction without creating technological fragmentation or unacceptable governance exposure.

Research Objectives

- To evaluate the strategic mechanisms through which SMEs deploy generative AI architectures to optimize business decision-making frameworks.
- To analyze and quantify the structural cost-reduction pathways achieved via GenAI across Operations, Human Resources, Customer Support, and Marketing.
- To design an optimized cost-benefit model establishing the financial break-even horizons for SME AI integrations.
- To identify and construct corporate governance strategies addressing data privacy, compliance, and technological dependencies unique to smaller organizational footprints.

Literature Review and Theoretical Framework

The study is anchored in the Resource-Based View (RBV) and Dynamic Capabilities perspectives. Barney's RBV argues that sustainable advantage arises from valuable, rare, inimitable, and non-substitutable resources. In the GenAI context, the source paper locates advantage not in widely accessible foundation models themselves, but in their contextual orchestration with proprietary data, specialized workflows, and employee capability. This combination can convert internal information into operational knowledge and therefore function as a firm-specific capability.

Teece, Pisano, and Shuen's Dynamic Capabilities framework complements this argument by emphasizing the ability to sense, seize, and transform in response to environmental change. The source study argues that SMEs may benefit from flatter structures and shorter feedback loops, allowing them to operationalize AI-enabled market signals and process changes quickly. Contemporary GenAI research cited in the source also associates generative systems with productivity gains and a shift in the economics of knowledge work.

A central distinction in the paper is between analytical AI and generative AI. Analytical systems are described as specialized, data-intensive tools focused on prediction, prescription, forecasting, and anomaly detection, whereas generative systems provide natural-language and multimodal interfaces with broader cross-functional scope. The original comparison is retained below without changing its supplied entries.

Table 1. Operational Comparison of Analytical AI vs. Generative AI in Corporate Infrastructures

Dimension	Analytical AI Era (Predictive/Prescriptive)	Generative AI Era (Constructive/Cognitive)
Capital Expenditure (CapEx)	High: Custom model training, data science teams, high-performance compute.	Low: Subscription models, tiered API access, plug-and-play integrations.
Data Requirements	Massive datasets; highly cleaned, structured, and labeled data repositories.	Flexible: Zero-shot/few-shot prompts utilizing unstructured natural language data.

Interface Modality	Monolithic outputs: Dashboards, CSV matrices, proprietary statistical visualizations.	Natural Language: Conversational interfaces, multi-modal outputs (text, code, image).
Functional Scope	Highly specialized: Linear regression, demand forecasting, anomaly detection.	Horizontal: Cross-functional content generation, reasoning, summarization, orchestration.

Research Methodology

The source document describes the study as a mixed-methods, multi-case investigation encompassing 45 SMEs across manufacturing, professional services, and retail. The analysis maps functional AI-integration pathways and combines qualitative interpretation of organizational use cases with the quantitative cost and performance evidence reported in the document. The present journal version does not add a sampling frame, respondent profile, questionnaire instrument, inferential test, or data-collection procedure because these details are not specified in the supplied synopsis.

The analytical approach has three components. First, functional use cases are examined across decision support, customer operations, marketing, administration, and software/data systems. Second, the study applies the source document's marginal-cost framework to explain the break-even logic of AI adoption. Third, the reported three-year cost projections are compared across functional domains to identify where savings emerge after implementation. Governance and implementation risks are interpreted alongside these economic outcomes.

Results And Discussion

GenAI and Business Decision-Making

The source study identifies three principal decision-support pathways. First, natural-language interfaces can reduce the technical barrier to business intelligence by translating managerial questions into database queries and narrative summaries. Second, generative systems can support real-time scenario planning by combining operating histories, price signals, logistics information, and other contextual inputs to compare alternative responses. Third, AI-assisted synthesis of public and internal information can strengthen customer and market intelligence, helping SMEs identify emerging needs, competitor activity, and potential market gaps.

These mechanisms are consistent with the RBV and Dynamic Capabilities framing: the model itself is broadly accessible, while advantage depends on how effectively the SME connects it to proprietary context, validates outputs, and embeds it in rapid decision cycles. GenAI therefore functions less as a stand-alone technology and more as an orchestration layer for organizational knowledge and action.

Structural Cost-Reduction Pathways

Administrative and operational workflows represent a direct automation opportunity. The source paper describes GenAI applications in drafting routine HR documents, screening information, onboarding support, invoice and expense processing, and extraction of structured information from unstructured inputs. In customer support, the reported evidence for service-oriented SMEs states that Retrieval-Augmented Generation deployments reduced first-response latency by 68% and deflected up to 74% of standard support queries. These figures are retained as reported in the supplied document.

Marketing is another major pathway because generative systems can produce copy and creative variations internally and personalize outreach using customer histories and behavioral information. The source also reports software-development efficiency gains of 40% to 60% for professional developers using GenAI accelerators, enabling faster delivery of internal tools, automation scripts, and integrations. Taken together, these pathways shift staff effort away from repetitive production and toward higher-value review, exception handling, and strategic work.

Financial Optimization and ROI Model

The source paper models traditional operating cost as $C_{\text{traditional}}(V) = F + \beta V$, where F is fixed overhead, V is transaction volume, and β is the marginal variable cost per transaction. Under AI adoption, the model becomes $CAI(V) = F + IAI + \mu + \gamma V$, where IAI represents setup and integration expenditure, μ represents recurring infrastructure and API consumption cost, and γ is the lower AI-enabled variable cost coefficient.

Equating the two cost functions yields the break-even volume $V^* = (IAI + \mu) / (\beta - \gamma)$. The framework therefore predicts that the required adoption volume declines as the gap between the traditional marginal cost β and the AI-enabled marginal cost γ increases. This model is conceptual in the supplied study; no firm-level coefficient estimates are provided, so the journal version does not infer numerical break-even volumes beyond the reported projections.

Table 2. Functional Cost-Benefit Projections Over 3-Year Lifecycle (Normalized Values in USD)

Functional Domain	Baseline Spend	Year 1 (Implementation)	Year 2 (Optimization)	Year 3 (Mature State)	Net Savings (%)
Customer Operations	120,000	135,000 (Incl. Setup)	65,000	40,000	66.6%
Marketing & Media	85,000	50,000	35,000	25,000	70.5%
Administrative Support	95,000	80,000	55,000	45,000	52.6%
Data/Software Systems	110,000	115,000	70,000	60,000	45.4%
Aggregated Total	410,000	380,000	225,000	170,000	58.5%

The reported projections show different adoption trajectories by function. Customer Operations rises above baseline during Year 1 because setup expenditure is included, then falls to USD 65,000 in Year 2 and USD 40,000 in the mature state, with reported net savings of 66.6%. Marketing & Media records the largest stated saving at 70.5%, while Administrative Support and Data/Software Systems report 52.6% and 45.4%, respectively. At the aggregate level, spending moves from USD 410,000 baseline to USD 380,000 in Year 1, USD 225,000 in Year 2, and USD 170,000 in Year 3, corresponding to the reported 58.5% net saving.

Risks, Governance, And Implementation

The economic case must be considered together with governance risk. The source study highlights data confidentiality as a primary concern when staff enter customer, financial, or proprietary information into public generative systems. It recommends enterprise-grade arrangements, zero-retention controls, and data filtering or anonymization before sensitive information leaves the organization. These safeguards are particularly important for SMEs because a single disclosure incident may have disproportionate legal and commercial consequences.

A second risk is hallucination and algorithmic bias. Because generative models can produce fluent but incorrect outputs, the source recommends a Human-in-the-Loop validation framework in which high-risk outputs—such as legal terms, vendor pricing, or safety-related material—receive mandatory domain-expert review. Lower-risk outputs may be subject to automated validation. A third concern is vendor lock-in; the proposed response is a flexible multi-model architecture that permits workloads to move between providers or open-source models as price, performance, and risk requirements change.

Strategic Implementation Roadmap

The original four-phase roadmap is retained in condensed journal form. Phase 1 (Weeks 1–4) conducts an operational audit, establishes acceptable-use policies, and secures enterprise access with appropriate data controls. Phase 2 (Weeks 5–12) pilots a high-impact, lower-risk use case and benchmarks performance against predefined KPIs. Phase 3 (Weeks 13–24) integrates validated systems with enterprise data through secure interfaces and expands staff training and quality-validation practices. Phase 4 is continuous: firms monitor spending and model performance, audit for drift or bottlenecks, and route suitable lower-complexity tasks to lower-cost models.

Conclusion

The study indicates that generative AI can alter the economics of digital capability for SMEs by making advanced analysis, content generation, customer support, and software assistance accessible through comparatively flexible interfaces. Within the evidence reported in the source document, the strongest business case combines faster decision support with reductions in recurring operating effort, while the three-year projections illustrate the possibility of substantial mature-state savings after implementation costs are absorbed.

The findings should nevertheless be interpreted within the information available in the supplied study. The paper states a 45-SME mixed-methods basis but does not provide the detailed sampling procedure, respondent characteristics, instrument, or firm-level statistical outputs needed for broader causal generalization. Accordingly, the results are best treated as multi-case evidence and strategic projections rather than population-level estimates. Future research should report the underlying sampling and measurement procedures, compare adoption outcomes by sector and firm size, and test realized savings against projected savings over longitudinal implementation periods.

For practitioners, the central implication is that GenAI adoption should be treated as an organizational capability rather than a collection of isolated tools. Sustainable value depends on selecting high-frequency use cases, connecting models to relevant proprietary context, validating outputs according to risk, protecting confidential data, and maintaining flexibility across model providers. Under these conditions, SMEs may convert accessible generative technologies into faster decisions and structurally lower operating costs.

References

1. Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120.
2. Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. National Bureau of Economic Research, Working Paper No. 31161.
3. Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., & Yee, L. (2023). The economic potential of generative AI: The next productivity frontier. McKinsey & Company Research Report.
4. Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533.
5. Google Workspace & Deloitte. (2024). Small Business, Big Tech: Accelerating SME Growth Through Accessible Artificial Intelligence. Deloitte Digital Insights Series.
6. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
7. Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (Working Paper No. 31161). National Bureau of Economic Research. <https://doi.org/10.3386/w31161>
8. Mariani, M., & Dwivedi, Y. K. (2024). Generative artificial intelligence in innovation management: A preview of future research developments. *Journal of Business Research*, 175, 114542. <https://doi.org/10.1016/j.jbusres.2024.114542>
9. Dell’Acqua, F., McFowland, E., III, Mollick, E., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science*, 37(2), 403–423. <https://doi.org/10.1287/orsc.2025.21838>