How Agentic AI is Reshaping Corporate Strategy: Opportunities and Challenges for Indian Businesses

Aakshi Jain
MBA Core, Narsee Monjee Institute of Management Studies, Mumbai, India

***Abstract*—Agentic artificial intelligence (AI) — systems that autonomously plan, decide, and execute multi-step tasks with limited human oversight — is emerging as the successor to generative AI in enterprise technology strategy. This paper examines how agentic AI is reshaping corporate strategy for Indian businesses, tracing its evolution from generative systems and applying seven established strategy frameworks — Porter's Five Forces, the Resource-Based View, VRIO, Dynamic Capability Theory, Value Chain Analysis, the McKinsey 7S Model, and the Business Model Canvas — to assess its strategic implications. Using a qualitative, secondary-data methodology grounded in consulting reports, company disclosures, and academic and practitioner literature, the study analyses ten firms across technology, retail, banking, and consulting — including Microsoft, Infosys, Tata Consultancy Services, Reliance Retail, Tata Digital, Flipkart, HDFC Bank, ICICI Bank, Accenture, and Deloitte — to identify how agentic AI is reconfiguring competitive positioning, value chains, and organisational capabilities. The findings indicate that agentic AI yields sustainable advantage only when combined with proprietary data, governance maturity, and workforce reskilling; absent these, the technology risks rapid commoditisation and competitive convergence. The paper concludes with managerial implications and a phased adoption framework for Indian firms, together with directions for future research.**

***Index Terms*—Agentic AI; Corporate Strategy; Digital Transformation; Dynamic Capabilities; Indian Businesses; Resource-Based View.**

I. INTRODUCTION

Between 2022 and 2024, generative artificial intelligence (AI) moved from research curiosity to boardroom priority, promising to automate drafting, summarising, and ideation across every business function. Yet by 2025 a gap had opened between adoption and impact: industry surveys found that a large majority of enterprises had deployed generative AI in at least one function, while most reported no measurable improvement in productivity, cost, or revenue as a result [1]. This "Generative AI Paradox" set the stage for a second wave of enterprise AI — agentic AI — in which systems no longer merely respond to prompts but autonomously set goals, plan multi-step tasks, invoke tools, and act with limited human oversight [2]. Analyst estimates place the global enterprise agentic AI market at roughly USD 5.3 billion in 2026, expanding toward USD 24.5 billion by 2030 at a compound annual growth rate above 46 percent [3], and enterprises with high-performing AI programmes are already committing more than a fifth of their digital budgets to AI-related infrastructure [2].

India sits at the centre of this transition, both as an adopter and as an economy whose services-export model is most exposed to it. Industry research indicates that roughly four in five Indian organisations are already exploring the development of autonomous AI agents [4], even as the country's technology, retail, banking, and consulting sectors race to translate that exploration into deployable strategy. The Government of India's National Strategy for Artificial Intelligence, first articulated by NITI Aayog in 2018 around the principle of "AI for All," anticipated much of this shift, framing AI as a driver of inclusive economic growth rather than a narrow productivity tool [5]. Six years on, agentic AI is testing whether that ambition can be reconciled with the competitive pressure Indian firms now face from both domestic rivals and global technology platforms.

Despite the volume of commentary on agentic AI, most of it is technology-centric — describing what agents can do — rather than strategy-centric — explaining what their diffusion means for how firms compete, organise, and capture value. This is a

significant omission for Indian businesses, which must decide, often with limited internal AI expertise, whether agentic AI is a genuine source of sustainable competitive advantage or a fast-commoditising capability that merely raises the cost of staying competitive. This paper addresses that gap by applying seven canonical strategy frameworks to the phenomenon of agentic AI and by comparing how ten technology, retail, banking, and consulting firms operating in India are translating agentic AI into strategic action.

The remainder of the paper proceeds as follows. Section II reviews the academic and practitioner literature on AI and strategy. Section III traces the evolution from generative to agentic AI. Section IV identifies the research gap this study addresses, and Section V sets out the research objectives, questions, and scope. Section VI describes the methodology. Section VII applies the seven strategy frameworks to agentic AI, and Section VIII presents the industry case studies. Sections IX to XII discuss the findings, managerial implications, limitations, and conclusions.

II. LITERATURE REVIEW

The strategic management literature has long used the Resource-Based View and Dynamic Capability Theory to explain how firms convert technology investment into competitive advantage, arguing that advantage flows not from technology adoption itself but from how firms combine technology with complementary, hard-to-imitate organisational resources. Recent practitioner research extends this logic to agentic AI, treating it not as a single tool but as a new category of organisational actor that firms must integrate into their operating model [6].

Agentic AI is most commonly defined in contrast to generative AI: where generative systems are reactive, producing content in response to a prompt, agentic systems are proactive, capable of setting goals, decomposing them into executable steps, invoking external tools and data sources, and adapting their plan as circumstances change, typically built around a reasoning core, a planning module, tool integration, and memory or feedback loops [1]. Analyst firms characterise 2026 as an inflection year for this shift: one major advisory firm projects that task-specific AI agents will be embedded in roughly 40 percent of enterprise applications by the end of 2026, up from under 5 percent in 2025, even as the same firm warns that most current agentic AI initiatives remain early-stage experiments driven more by hype than by proven business cases [7].

A joint 2025–2026 study by MIT Sloan Management Review and Boston Consulting Group frames the resulting management challenge around four tensions that agentic AI introduces into the enterprise — over control, investment, governance, and organisational design — concluding that enthusiasm for agentic AI is running ahead of most organisations' readiness to manage it [6]. This concern about readiness converges with a second stream of literature warning about the erosion of competitive advantage itself: Harvard Business Review researchers describe an “agentic convergence trap,” in which firms that deploy AI systems trained on similar market data and optimised toward similar objectives arrive at near-identical strategic decisions, eroding the very differentiation that strategy is meant to protect [8]. Boston Consulting Group's global executive survey work adds an investment-allocation dimension: although a large share of chief executives already use AI agents personally on a weekly basis, the great majority of enterprise AI investment still targets operational productivity rather than the strategic decision-making processes where BCG argues the next competitive advantage will be won [9]. A companion BCG study of corporate strategy functions similarly finds that AI tools have so far delivered consistently positive impact mainly in market intelligence and research, with their fuller strategic potential — reshaping how strategy itself is formulated — still largely untapped [10].

India's own policy literature anticipates several of these tensions. NITI Aayog's “Responsible AI for All” approach document, developed with the World Economic Forum as knowledge partner, sets out principles of safety, transparency, accountability, and inclusion intended to guide AI deployment across sectors [11], while its 2025 report on AI for inclusive societal development extends this agenda to India's vast informal workforce, arguing that AI should be designed to amplify rather than displace the skills of workers outside the formal economy [12]. Read together, this literature suggests that the strategic promise of agentic AI for Indian firms cannot be assessed on technological grounds alone; it must be evaluated against governance readiness, competitive convergence risk, and the distributive consequences of

automation — precisely the gap this paper sets out to address.

III. EVOLUTION OF AI: FROM GENERATIVE TO AGENTIC INTELLIGENCE

Generative AI systems, popularised from 2022 onward, are fundamentally reactive: given a prompt, they produce text, code, images, or analysis, but they neither set their own objectives nor act on the world without a human initiating and directing each step. Agentic AI systems differ along four dimensions: they can accept a goal rather than a single instruction; they plan a sequence of sub-tasks needed to reach that goal; they call external tools, application programming interfaces, and data sources to execute those sub-tasks; and they retain memory of outcomes to adapt subsequent actions, operating with progressively less human oversight as trust in the system grows [1]. Table I summarises this distinction.

TABLE I GENERATIVE AI VERSUS AGENTIC AI

Dimension	Generative AI	Agentic AI
Core function	Produces content on request	Plans & executes multi-step goals
Autonomy	Low; needs a prompt each step	High; acts with limited oversight
Interaction pattern	Ask-and-answer	Observe-and-act
Primary value driver	Individual productivity	End-to-end process automation
Illustrative example	Chatbot answering a query	Agent resolving a claim end-to-end

This shift is already visible in enterprise software. Vendors are rebuilding core productivity and enterprise-resource-planning suites around autonomous “agentic ERP” and multi-agent orchestration rather than single-purpose copilots, extending governance layers to manage agent identity, permissions, and audit trails as agents are given the ability to take real actions rather than merely draft suggestions for a human to approve. Advisory firms nonetheless caution that the technology is entering a period of recalibration: agentic AI is expected to sit at the “peak of inflated expectations” on at least one major hype-cycle model before moving into a “trough of disillusionment” through 2026, even as measured adoption of task-specific agents continues to accelerate in parallel [7]. For corporate strategists, the practical implication is that agentic AI should be evaluated less as a single technology purchase and more as a capability platform whose competitive value depends on how it is embedded into a firm's specific processes and data — the theoretical question taken up in Section VII.

IV. RESEARCH GAP

Four gaps motivate this study. First, the bulk of agentic AI literature is produced by, and about, North American and European enterprises; systematic strategic analysis specific to Indian firms — with their distinctive mix of IT-services export dependence, large domestic conglomerates, and a fast-evolving AI policy environment — remains scarce. Second, canonical strategy frameworks such as Porter's Five Forces and VRIO are rarely applied explicitly to agentic AI; most commentary substitutes general enthusiasm or alarm for structured strategic reasoning. Third, existing company-level accounts of agentic AI adoption are largely descriptive, drawn from press releases and product announcements, without a comparative lens that allows patterns to be identified across sectors. Fourth, governance and workforce risk are typically discussed in isolation from corporate strategy, even though — as Sections VII to IX argue — they are central determinants of whether agentic AI translates into sustainable advantage. This paper addresses all four gaps through a framework-based, multi-sector, India-focused comparative analysis.

V. RESEARCH OBJECTIVES, QUESTIONS, AND SCOPE

A. Objectives

This study pursues five objectives: (1) to understand the evolution from generative AI to agentic AI; (2) to analyse how agentic AI influences corporate strategy; (3) to identify strategic opportunities for Indian firms; (4) to examine the organisational and governance challenges of implementation; and (5) to propose a strategic framework for successful adoption.

B. Research Questions

The analysis is guided by five questions: (1) How is agentic AI changing strategic decision-making? (2) Can agentic AI become a source of sustainable

competitive advantage? (3) Which industries in India are most likely to benefit? (4) What organisational capabilities are required to capture that benefit? (5) What risks should Indian business leaders anticipate?

C. Scope and Delimitations

The study is delimited to four sectors — technology and IT services, retail and e-commerce, banking, and consulting — and to ten firms selected purposively within them: Microsoft, Infosys, and Tata Consultancy Services (technology); Reliance Retail, Tata Digital, and Flipkart (retail); HDFC Bank and ICICI Bank (banking); and Accenture and Deloitte (consulting). These firms were chosen because each has made a public, verifiable strategic commitment to agentic AI between 2025 and 2026, allowing evidence-based comparison. The analysis reflects publicly available information current to July 2026 and does not extend to primary financial or operational data that firms have not disclosed.

VI. METHODOLOGY

This is a qualitative, exploratory study based entirely on secondary data, appropriate given how recently agentic AI has entered mainstream corporate strategy and the resulting absence of longitudinal primary data. Evidence was drawn from three categories of source: (a) reports and publications from global consulting and advisory firms, including McKinsey, Boston Consulting Group, Deloitte, Accenture, Bain, and Gartner; (b) practitioner and academic outlets, including Harvard Business Review and MIT Sloan Management Review; and (c) primary organisational disclosures, including company newsroom releases, investor and strategy communications, and Government of India policy documents issued by NITI Aayog.

Company selection followed a purposive sampling logic designed to maximise sectoral spread while holding data availability constant: within each of four sectors most exposed to agentic AI in India, firms were included only if they had made a specific, named, and independently verifiable agentic AI initiative public by July 2026.

The seven strategy frameworks in Section VII were applied qualitatively rather than through numerical scoring: each framework was used as an interpretive lens to organise and synthesise the evidence gathered on agentic AI's strategic implications and on the ten case-study firms, consistent with how these frameworks are conventionally used in exploratory strategic analysis. Findings were triangulated across at least two independent sources wherever possible to reduce the risk of single-source bias inherent in press-release-driven evidence. The absence of interview or survey-based primary data is acknowledged as a limitation and addressed further in Section XI.

VII. STRATEGIC ANALYSIS OF AGENTIC AI

A. Porter's Five Forces

Agentic AI is reshaping every one of Porter's five competitive forces. The threat of new entrants is rising: cloud-based agent-orchestration tooling and increasingly capable foundation models lower the technical barrier for AI-native start-ups to challenge incumbents with a fraction of their headcount, a dynamic strategy researchers describe as placing incumbents' traditional scale advantages under direct threat [13]. The bargaining power of buyers is also increasing, as emerging agentic-commerce protocols allow AI shopping agents to compare prices and negotiate on a buyer's behalf across merchants; one such protocol, co-developed by Google and Shopify, already counts Flipkart among more than twenty retail and payments endorsers, alongside Walmart, Target, Visa, and Mastercard [14]. The bargaining power of suppliers is shifting toward the small number of hyperscale cloud and foundation-model providers on which most enterprise agent stacks depend. The threat of substitutes is rising as agent-mediated interfaces threaten to substitute traditional websites, applications, and human-staffed service channels. Finally, competitive rivalry is intensifying in a specific way: because many firms are licensing similar third-party agent platforms trained on similar data, agentic capability risks becoming a competitive-parity requirement rather than a differentiator — the convergence risk identified in Section II [8].

B. VRIO Framework

Applying the VRIO framework separates the parts of agentic AI that create advantage from the parts that merely create parity. Table II summarises the analysis: a generic, off-the-shelf agent platform is valuable but neither rare nor difficult to imitate, so on its own it yields only competitive parity. A proprietary,

workflow-tuned agentic layer built on a firm's own data — such as the composable "agentic services" platforms Indian IT majors are building — is valuable and, at least initially, rare, though only partially inimitable once competitors observe and replicate the approach. Governance and agent-orchestration capability — the organisational ability to deploy, permission, audit, and continuously improve fleets of agents safely — is the resource most likely to be valuable, rare, and difficult to imitate because it is path-dependent, built through accumulated organisational learning rather than purchased outright. The implication is that the underlying agentic AI technology is commoditising rapidly, while sustainable advantage is migrating toward the proprietary data and governance capability that surround it.

TABLE II VRIO ANALYSIS OF AGENTIC AI AS A STRATEGIC RESOURCE

Resource	V	R	I	O	Outcome
Generic / off-the-shelf agent platform	Y	N	N	Varies	Competitive parity
Proprietary, workflow-tuned agentic layer	Y	Y*	Partial	Firm-dep.	Temporary advantage
Governance & orchestration capability	Y	Y	Y*	If invested	Sustained advantage

* Rarity and inimitability are path-dependent and erode over time as competitors observe and replicate governance practices; V = Valuable, R = Rare, I = Inimitable, O = Organised.

C. Resource-Based View

Extending VRIO, the Resource-Based View suggests that agentic AI should be understood as a bundle of complementary resources — proprietary data, domain-specific workflow knowledge, governance infrastructure, and specialised talent — rather than as a single resource in its own right. Indian IT-services firms illustrate this logic directly: rather than reselling third-party agent technology unmodified, Infosys has built a composable, model-agnostic "agentic services" layer that sits above partner models from Anthropic, OpenAI, and others [15], while Tata Consultancy Services has built a platform-led orchestration layer that aggregates and governs multiple generative and agentic AI services in a single interface [16]. In both cases the strategic asset being built is not the underlying model — which the firm does not own — but the orchestration, governance, and integration layer wrapped around it, precisely the kind of valuable, rare, and imperfectly imitable resource bundle the Resource-Based View predicts will sustain advantage.

D. Dynamic Capability Theory

Dynamic Capability Theory frames agentic AI adoption as a three-stage process of sensing, seizing, and transforming. Sensing involves identifying which internal processes are suited to agentic automation; seizing involves rapid piloting and scaling of the highest-value opportunities; and transforming involves reconfiguring organisational structure, governance, and workforce around the resulting capability. Evidence from enterprise technology leaders suggests only a small minority of firms — on the order of one in twenty — have reached genuine "agent-first" maturity, and that the firms furthest ahead prioritise speed of deployment first, growth second, and cost reduction third, achieving cost reductions on the order of 15 to 20 percent across functions as a result [17]. Critically, the same research finds that roughly 70 percent of what determines success in agentic AI implementation is organisational change management rather than algorithmic sophistication [17] — a finding echoed in the transforming undertaken by consulting firms themselves, discussed further in Section VIII-D.

E. Value Chain Analysis

Agentic AI applications are already visible across the primary and support activities of Porter's value chain, summarised in Table III. In procurement, autonomous sourcing agents can evaluate and negotiate with vendors directly; in operations, enterprise software vendors are redesigning core planning and resource-management systems around agentic workflows; in marketing, multi-agent systems can plan and execute campaigns end-to-end, though survey evidence suggests fewer than one in ten marketing organisations currently run campaigns with multiple agents operating fully autonomously [9]; in sales, customer-relationship-management copilots increasingly qualify and route leads with minimal human input; in human resources, agents support onboarding and reskilling at scale; and in customer service, banks

deploying agentic AI in know-your-customer and anti-money-laundering workflows have reported productivity gains ranging from 200 to more than 2,000 percent [7].

TABLE III AGENTIC AI ACROSS THE VALUE CHAIN

Value Chain Activity	Illustrative Agentic AI Application
Procurement	Autonomous sourcing & vendor-negotiation agents
Operations	Agentic ERP & self-monitoring supply chains
Marketing	Multi-agent campaign planning & execution
Sales	CRM copilots that qualify and route leads
Human Resources	Onboarding & reskilling agents
Customer Service	Autonomous KYC/AML & claims-resolution agents

F. McKinsey 7S Framework

The McKinsey 7S model highlights the breadth of organisational change agentic AI requires, beyond strategy alone. Structurally, firms are creating new agent-orchestration and AI-governance functions and flattening escalation paths as agents absorb routine, first-line tasks. In systems, firms must build agent identity, permissioning, and audit-trail infrastructure comparable to the governance control-plane approach large technology vendors are now offering enterprise customers [18]. In shared values, Indian firms have a ready-made anchor in NITI Aayog's responsible-AI principles of safety, transparency, and accountability [11]. In skills, structured reskilling programmes — such as one Indian bank's generative-AI academy, which has already trained tens of thousands of staff [19] — are emerging as a precondition for successful deployment rather than an afterthought. In style, leadership language is shifting from treating AI as a tool employees operate to treating it as a "teammate" or digital colleague that must be managed, a reframing strategy researchers argue is itself a leadership capability [6]. Finally, in staff, Section IX shows that this shift is beginning to reshape the composition, not merely the skills, of the Indian workforce.

G. Business Model Canvas

Applying the Business Model Canvas, agentic AI is altering revenue streams, cost structure, customer relationships, and key activities simultaneously. On revenue, IT-services firms are shifting from time-and-materials billing toward "services-as-software" models priced on outcomes delivered by agent-driven platforms [20]. On cost structure, agent-first firms report the 15-to-20-percent efficiency gains noted in Section VII-D, but transformation is not free: one global consulting firm disclosed roughly USD 865 million in one-time restructuring charges as part of an AI-driven reskilling and reorganisation programme [21]. On customer relationships, agentic AI is shifting an increasing share of interactions from human-mediated to agent-mediated channels, from banking chatbots to autonomous shopping assistants. On key activities, the centre of value creation is moving from application development toward agent orchestration, workflow redesign, and governance — the same shift Section VII-B identifies as the true locus of sustainable advantage.

VIII. INDUSTRY CASE STUDIES

Table IV summarises the ten firms analysed before the sector-by-sector discussion that follows.

TABLE IV CASE-STUDY FIRMS AT A GLANCE

Company	Sector	Signature Agentic AI Move	Strategic Posture
Microsoft	Technology	Agent 365 governance control plane	Platform & governance leadership
Infosys	IT Services	Topaz Fabric agentic services suite	Neutral multi-model integrator
Tata Consultancy Services	IT Services	AI WisdomNext + AI factories	Platform-led infrastructure play
Reliance Retail	Retail	Reliance Intelligence; Jio Call Agent	Vertically integrated infrastructure bet
Tata Digital	Retail	Tata Neu, powered by TCS	Federated, partnership-led
Flipkart	E-commerce	"Flippi" assistant; UCP endorsement	Open standards participation
HDFC Bank	Banking	Neev platform; GenAI Academy	AI-first, unified governance
ICICI Bank	Banking	iPal evolution; Innovation Labs	Staged, risk-managed diffusion

Company	Sector	Signature Agentic AI Move	Strategic Posture
Accenture	Consulting	AI Refinery; reinvention programme	Vendor + internal proof-of-concept
Deloitte	Consulting	State of AI in Enterprise research	Research-led advisory positioning

A. Technology and IT Services: Microsoft, Infosys, and Tata Consultancy Services

Microsoft is positioning itself as the governance layer for enterprise agentic AI rather than solely a model provider: its Agent 365 platform, made generally available in May 2026, functions as a control plane that lets enterprises discover, govern, and secure agents built on Microsoft's own tools as well as external and locally run agents, extending network-level controls to detect unsanctioned "shadow AI" usage [18], while agentic capability is simultaneously being built directly into everyday productivity applications [22]. Infosys has pursued a platform-and-partnership strategy through Infosys Topaz Fabric, a composable, model-agnostic agentic-services suite through which it has struck collaborations with Anthropic, OpenAI, Intel, and other AI providers, positioning itself to capture a share of the roughly USD 300–400 billion AI-services opportunity that a joint Nasscom–McKinsey study projects for the sector by 2030 [20], [15]. Tata Consultancy Services has taken a similar platform-led route through TCS AI WisdomNext, an orchestration layer that aggregates multiple generative and agentic AI services with embedded governance, which analysts at Everest Group named a market leader in AI and generative AI services in March 2026, citing TCS's co-innovation with hyperscalers and price competitiveness as differentiators [23], [16]. Across all three firms, the common strategic pattern is repositioning from linear, labour-arbitrage-based service delivery toward platform-led, intellectual-property-owning models designed to defend margins against the disintermediation agentic AI threatens.

B. Retail and E-Commerce: Reliance Retail, Tata Digital, and Flipkart

Reliance has pursued the most capital-intensive strategy of the three, pledging roughly USD 110 billion toward AI infrastructure, establishing a dedicated subsidiary, Reliance Intelligence, to build gigawatt-scale AI data centres, and forming AI joint ventures with Google and Meta [24]; at its 2026 annual general meeting the group unveiled Jio Call Agent, a voice-based AI assistant able to join phone calls to transcribe, summarise, and complete tasks such as booking a cab or a restaurant reservation for its more than 500 million telecom subscribers [25]. This vertically integrated, infrastructure-heavy approach tightly couples Reliance's telecom, retail, and AI ambitions. Tata Digital has taken a more federated route: its Tata Neu "super-app" integrates group brands such as BigBasket, Croma, and 1mg, but its AI capability is delivered primarily through group technology arm Tata Consultancy Services rather than through a unified, Reliance-style proprietary data backend, reflecting a more partnership-driven, less centrally orchestrated strategy. Flipkart, majority-owned by Walmart, has pursued interoperability rather than a walled garden: alongside launching its own conversational shopping assistant, "Flippi," it is among the retailers exploring agentic commerce that have endorsed Google's Universal Commerce Protocol for agent-mediated shopping, alongside Walmart, Target, and major payments networks [14] — a strategic bet that participating in open standards will protect market access as AI shopping agents proliferate [26]. The sector-level pattern is a divergence between infrastructure-led vertical integration and standards-based partnership, reflecting each firm's distinct resource base and ownership structure.

C. Banking: HDFC Bank and ICICI Bank

HDFC Bank has articulated an explicit "AI-first" ambition, targeting AI involvement in the large majority of customer interactions, anchored by its Neev generative AI platform and a GenAI Academy that has already trained tens of thousands of staff, with every model tied to a single customer identity and governance layer [19]. ICICI Bank has favoured a more staged, risk-managed diffusion path: its AI-driven assistant, iPal, has evolved over several product generations, and the bank has used its Innovation Labs to pilot generative AI tools such as an internal business-requirement-specification generator before extending similar capability to customer-facing channels [27]. Sector-wide, the productivity case for agentic AI in banking is substantial — as noted in

Section VII-E, banks applying agentic AI to know-your-customer and anti-money-laundering workflows report some of the largest productivity gains observed in any industry — but deployment remains bounded by regulatory caution, making staged piloting and unified governance the true differentiators rather than raw AI capability itself.

D. Consulting: Accenture and Deloitte

Consulting firms occupy a distinctive position in the agentic AI economy, simultaneously selling transformation and undergoing it internally. Accenture has booked roughly USD 5 billion in AI-related revenue since the launch of ChatGPT in 2022, expanded its AI Refinery ecosystem, and taken venture stakes in agentic-AI specialists such as Profitmind, for retail pricing and inventory decisions, and Netomi, for agentic customer experience [28]; it has paired this external push with an internal "reinvention" programme carrying roughly USD 865 million in restructuring charges, explicitly prioritising upskilling first and workforce exit only where reskilling is not viable [21]. Deloitte has positioned itself primarily as researcher and governance advisor: its 2026 State of AI in the Enterprise survey of more than 3,000 leaders across 24 countries found that roughly three-quarters of organisations expect at least moderate agentic AI use within two years, yet only about one in five currently has a mature agent-governance model in place [29] — a gap its own agentic AI strategy guidance for clients is explicitly designed to close [30]. The shared pattern across both firms is that credibility in advising clients on agentic AI increasingly depends on firms having already reorganised themselves around it.

IX. DISCUSSION

Read together, the frameworks in Section VII and the cases in Section VIII point to a consistent conclusion: agentic AI is bifurcating firms into a small set of agent-first leaders and a much larger group of followers, and in India that bifurcation cuts across firm type rather than simply firm size. Capital-rich conglomerates such as Reliance can pursue infrastructure-heavy, vertically integrated strategies that smaller rivals cannot replicate; IT-services majors such as TCS and Infosys are instead monetising orchestration and governance intellectual property precisely because the underlying models are not proprietary to them; and regulated banks such as HDFC and ICICI are constrained to staged, governance-first deployment regardless of their appetite for speed.

The convergence risk identified in the literature review is visible in the case evidence too. Where multiple Indian and multinational firms license similar third-party agent platforms — the same underlying foundation models, the same commercially available orchestration tooling — the strategic differentiator observed across every case in Section VIII is not the technology itself but what each firm has built around it: Infosys's and TCS's proprietary orchestration layers, HDFC's unified governance architecture, Accenture's internally proven reinvention playbook. This is precisely the pattern VRIO analysis in Section VII-B predicts: technology commoditises, while the organisational capability wrapped around it does not.

The clearest tension specific to India lies in the labour market. Agentic AI offers Indian IT-services firms a way to defend against commoditisation of the coding, testing, and business-process work that built the sector, but it simultaneously threatens the labour-arbitrage model that made that sector globally competitive in the first place. Even as the largest Indian IT firms have continued to commit to hiring tens of thousands of fresh graduates [31], the same firms' combined headcount fell by tens of thousands of roles in the 2026 financial year, concentrated in traditional, process-driven service roles rather than AI, cloud, or platform-engineering positions [31]. Independent estimates suggest that more than 60 percent of India's formal-sector service jobs could face significant automation exposure by 2030 [32], producing what analysts describe as an "hourglass" labour market in which senior and specialised roles remain in demand even as the entry-level rung that has historically absorbed India's engineering graduates narrows [32].

Finally, governance remains the binding constraint on realising agentic AI's strategic promise almost everywhere. Global surveys converge on a similar figure: only around one in five enterprises has a mature governance model for autonomous agents even as most expect at least moderate deployment within two years [29], and at least one major advisory firm forecasts that more than 40 percent of agentic AI projects will be cancelled by 2027 because of governance gaps, unclear business value, and

inadequate risk controls rather than any shortfall in the underlying technology [33], [34]. For Indian firms with comparatively less mature enterprise AI governance infrastructure than global peers, this constraint is likely to bind earlier rather than later.

X. MANAGERIAL IMPLICATIONS

For Indian business leaders, the analysis in Sections VII to IX yields six practical implications:

1) Treat agentic AI primarily as an organisational-capability investment rather than a software purchase, since the evidence in Section VII-D indicates that roughly 70 percent of implementation success depends on change management rather than algorithmic sophistication.

2) Differentiate through proprietary data, workflow integration, and governance rather than through the underlying agent technology itself, which Section VII-B shows is commoditising rapidly across competitors.

3) Build tiered, risk-scoped governance rather than a single uniform policy for all agents, to avoid the project-cancellation risk documented in Section IX.

4) Pair every agent deployment with a structured reskilling programme rather than reactive workforce reduction, following the upskilling-first approach observed at both HDFC Bank and Accenture, to protect institutional knowledge and the entry-level talent pipeline documented in Section IX.

5) Engage early with India's emerging AI-governance architecture — NITI Aayog's responsible-AI principles and sector-specific frameworks under development for financial services — rather than treating regulatory alignment as a late-stage compliance exercise.

6) Sequence adoption using the sense–seize–transform logic of Dynamic Capability Theory: pilot narrowly, scale the two or three highest-value processes quickly, and only then redesign structure and roles around the resulting capability, rather than pursuing diffuse, unfocused experimentation across the organisation simultaneously.

XI. LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

This study has four principal limitations. It relies entirely on secondary, publicly available data rather than primary interviews or surveys, which limits insight into firms' internal decision-making. The evidence base is evolving rapidly, and specific facts and figures cited here may be superseded within months given the pace of change documented throughout the paper. The strategy frameworks were applied qualitatively rather than through a validated quantitative scoring instrument, which aids interpretability but limits statistical comparability across firms. Finally, the purposive, ten-firm sample — while spanning four sectors — is not randomly drawn and cannot support claims of generalisability to the wider population of Indian businesses, particularly small and medium enterprises, which are not represented in this study.

Four directions for future research follow naturally. First, primary survey or interview-based research is needed to assess agentic AI governance maturity directly within Indian firms rather than inferring it from public disclosures. Second, longitudinal research should track whether early Indian agentic AI adopters sustain a competitive advantage over time or converge with followers, providing an empirical test of the convergence-trap thesis discussed in Section II within the Indian context specifically. Third, sector-specific studies — particularly on the labour-market restructuring of India's IT-services sector and on sector-specific financial-services governance frameworks as they mature — would extend the discussion in Section IX considerably. Fourth, future work should develop and validate quantitative instruments for scoring organisational AI capability against VRIO and Resource-Based View criteria, enabling more rigorous cross-firm and cross-country comparison than the qualitative approach used here.

XII. CONCLUSION

Agentic AI is shifting the central question of corporate AI strategy from “which AI tools should we adopt” to “how should our organisation be redesigned around autonomous digital actors.” For Indian businesses, this analysis suggests the opportunity is real but conditional: applying Porter's Five Forces, VRIO, the Resource-Based View, Dynamic Capability Theory, Value Chain Analysis, the McKinsey 7S Model, and the Business Model Canvas to evidence from ten technology, retail, banking, and consulting firms shows consistently that sustainable advantage accrues not to whichever firm adopts agentic AI first, but to

the firm that pairs it with proprietary data, disciplined governance, and genuine workforce transformation. Firms that treat agentic AI as a plug-in technology purchase are likely to find their advantage converge away within a competitive cycle or two; firms that treat it as a multi-year organisational-capability build are better positioned to capture and defend the gains this paper documents. This structured, framework-based, comparative assessment is offered as a foundation for the more rigorous empirical research on Indian agentic AI strategy that Section XI argues is now needed.

REFERENCES

[1] Kore.ai, "What is Agentic AI? Meaning, Use Cases & How It Works," Kore.ai Blog, 2026. [Online]. Available: https://www.kore.ai/blog/what-is-agentic-ai

[2] McKinsey & Company, "Reimagining tech infrastructure for agentic AI," McKinsey Insights, 2026. [Online]. Available: https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/reimagining-tech-infrastructure-for-and-with-agentic-ai

[3] Grand View Research, "Enterprise Agentic AI Market Size & Share Report, 2025-2030," 2026. [Online]. Available: https://www.grandviewresearch.com/industry-analysis/enterprise-agentic-ai-market-report

[4] Deloitte India, "Driving Adoption of the 'Human with Agentic AI' Era," Deloitte Insights, 2026. [Online]. Available: https://www.deloitte.com/in/en/services/consulting/services/engineering-ai-data/turning-agentic-ai-into-business-value.html

[5] NITI Aayog, National Strategy for Artificial Intelligence, Government of India, New Delhi, India, 2018.

[6] MIT Sloan Management Review and Boston Consulting Group, "The Emerging Agentic Enterprise: How Leaders Must Navigate a New Age of AI," MIT Sloan Management Review, 2025. [Online]. Available: https://sloanreview.mit.edu/projects/the-emerging-agentic-enterprise-how-leaders-must-navigate-a-new-age-of-ai/

[7] Generative, Inc., "Agentic AI in 2026: How AI Evolved from Chatbots to Autonomous Agents," 2026. [Online]. Available: https://www.generative.inc/agentic-ai-in-2026-how-ai-went-from-chatting-to-doing

[8] P. van Esch, Y. G. Cui, and J. S. Black, "Beware the Agentic Convergence Trap," Harvard Business Review, May 2026. [Online]. Available: https://hbr.org/2026/05/beware-the-agentic-convergence-trap

[9] Boston Consulting Group, "How AI Decision Agents Transform Strategy," BCG Publications, 2026. [Online]. Available: https://www.bcg.com/publications/2026/how-ai-decision-agents-transform-strategy

[10] Boston Consulting Group, "The Corporate Strategy Function in an AI-First World," BCG Publications, 2026. [Online]. Available: https://www.bcg.com/publications/2026/the-corporate-strategy-function-in-an-ai-first-world

[11] NITI Aayog, Responsible AI for All: Approach Document for India, Part 1 - Principles for Responsible AI, Government of India, New Delhi, India, Feb. 2021.

[12] NITI Aayog, AI for Inclusive Societal Development, Government of India, New Delhi, India, Oct. 2025.

[13] Harvard Business Review, "How Agentic AI Supercharges Startups and Threatens Incumbents," Jul. 2026. [Online]. Available: https://hbr.org/2026/07/how-agentic-ai-supercharges-startups-and-threatens-incumbents

[14] Opascope, "AI Shopping Assistant Guide 2026: Agentic Commerce Protocols," Apr. 2026. [Online]. Available: https://opascope.com/insights/ai-shopping-assistant-guide-2026-agentic-commerce-protocols/

[15] Infosys Ltd., "Infosys and Anthropic Announce Collaboration to Unlock AI Value across Complex, Regulated Industries," Infosys Newsroom, Feb. 2026. [Online]. Available: https://www.infosys.com/newsroom/press-releases/2026/advanced-enterprise-ai-solutions-industries.html

[16] Tata Consultancy Services, "TCS AI WisdomNext," TCS Solutions, 2026. [Online]. Available: https://www.tcs.com/what-we-do/services/artificial-intelligence/solution/enterprise-generative-ai-adoption-wisdomnext

[17] Boston Consulting Group, "Agentic AI Strategy for CIOs and CTOs in 2026," BCG Publications, 2026. [Online]. Available: https://www.bcg.com/publications/2026/agentic-ai-strategy-cio-cto

[18] Microsoft Corporation, "Microsoft Agent 365, Now Generally Available, Expands Capabilities and Integrations," Microsoft Security Blog, May 2026. [Online]. Available: https://www.microsoft.com/en-us/security/blog/2026/05/01/microsoft-agent-365-now-generally-available-expands-capabilities-and-integrations/

[19] Analytics Vidhya, "HDFC Bank's AI-First Revolution: A Case Study," Aug. 2025. [Online]. Available: https://www.analyticsvidhya.com/blog/2025/08/hdfc-bank-a-case-study/

[20] Infosys Ltd., "Infosys Unveils AI First Value Framework," Infosys Newsroom, Feb. 2026. [Online]. Available: https://www.infosys.com/newsroom/press-releases/2026/unveils-ai-first-value-framework.html

[21] AOL/Fortune, "Accenture's $865 Million Reinvention Includes Saying Goodbye to People Without the Right AI Skills," 2026. [Online]. Available: https://www.aol.com/articles/accenture-865-million-reinvention-includes-124500249.html

[22] Microsoft Corporation, "Powering Frontier Transformation with Copilot and Agents," Microsoft 365 Blog, Mar. 2026. [Online]. Available: https://www.microsoft.com/en-us/microsoft-365/blog/2026/03/09/powering-frontier-transformation-with-copilot-and-agents/

[23] Tata Consultancy Services, "TCS Named a Leader in Artificial Intelligence and Generative AI Services by Everest Group," TCS Newsroom, Mar. 2026. [Online]. Available: https://www.tcs.com/who-we-are/newsroom/analyst-reports/tcs-named-leader-artificial-intelligence-gen-ai-services-everest-group

[24] TechCrunch, "Reliance Unveils $110B AI Investment Plan as India Ramps Up Tech Ambitions," Feb. 2026. [Online]. Available: https://techcrunch.com/2026/02/19/reliance-unveils-110b-ai-investment-plan-as-india-ramps-up-tech-ambitions/

[25] TechCrunch, "Billionaire Ambani Wants AI in Every Call, App, and Home," Jun. 2026. [Online]. Available: https://techcrunch.com/2026/06/19/billionaire-ambani-wants-ai-in-every-call-app-and-home/

[26] Bain & Company, "Agentic AI in Retail: How Autonomous Shopping Is Redefining the Customer Journey," Bain Insights, May 2026. [Online]. Available: https://www.bain.com/insights/agentic-ai-in-retail-how-autonomous-shopping-redefining-customer-journey/

[27] Klover.ai, "ICICI Bank's AI Strategy: Analysis of Dominance in Banking," Jul. 2025. [Online]. Available: https://www.klover.ai/icici-bank-ai-strategy-analysis-of-dominance-in-banking/

[28] HFS Research, "The Big, Bold Accenture Reinvention," May 2026. [Online]. Available: https://www.hfsresearch.com/research/big-bold-accenture-reinvention/

[29] Deloitte, "The State of AI in the Enterprise, 2026 AI Report," Deloitte Insights, 2026. [Online]. Available: https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html

[30] Deloitte Insights, "Agentic AI Strategy," Deloitte Tech Trends 2026, Dec. 2025. [Online]. Available: https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2026/agentic-ai-strategy.html

[31] Zee Business, "IT Jobs: 1.34 Lakh Global Tech Layoffs, Yet TCS, Infosys and Cognizant Plan 80,000 Hires - Why?," Jun. 2026. [Online]. Available: https://www.zeebiz.com/companies/news-it-jobs-134-lakh-global-tech-layoffs-yet-tcs-infosys-and-cognizant-plan-80000-hires-why-396636

[32] Outlook India, "Slow Death of IT Sector Jobs: How AI Layoffs Are Reshaping India's Tech Workforce and Middle Class Dreams," Jul. 2026. [Online]. Available: https://www.outlookindia.com/national/slow-death-of-it-sector-jobs-how-ai-layoffs-are-

reshaping-indias-tech-workforce-and-middle-class-dreams

[33] Gartner, Inc., "Gartner Says Applying Uniform Governance Across AI Agents Will Lead to Enterprise AI Agent Failure," Gartner Newsroom, May 2026. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2026-05-26-gartner-says-applying-uniform-governance-across-ai-agents-will-lead-to-enterprise-ai-agent-failure

[34] R. Szczerba, "Why 40% of Agentic AI Projects May Be Canceled by 2027," Forbes, Jul. 2026. [Online]. Available: https://www.forbes.com/sites/robertszczerba/2026/07/07/why-40-of-agentic-ai-projects-may-be-canceled-by-2027/