



Role of Artificial Intelligence in E-Commerce: Evidence-Supported Impacts on Consumer Experience, Operations, and Risk in Global and Indian Markets

Miss. Bharti Pandurang Suvarnkar

Assistant Professor, Department of Commerce
Yeshwant Mahavidyalaya, Nanded.

Abstract

Artificial Intelligence (AI) has shifted e-commerce competition from scale-and-supply to prediction-and-orchestration: who anticipates demand, reduces friction, protects trust, and personalizes value at the lowest marginal cost. This paper synthesizes secondary evidence from multilateral datasets, policy documents, market reports, and peer-reviewed research to analyze how AI changes measurable outcomes across the e-commerce value chain. Using a structured evidence-mapping method, the study integrates global indicators of e-commerce scale (e.g., UNCTAD estimates of total business e-commerce sales), national market trajectories in India, and operational performance proxies such as conversion frictions and fraud loss projections. Four analytical tables are used to connect AI capabilities (ML, NLP, computer vision, and generative AI) to e-commerce functions, quantify adoption-and-scale signals, compare pre-AI versus AI-enabled performance patterns, and evaluate risk-control trade-offs under emerging regulatory constraints. Results indicate that

AI-enabled personalization is consistently associated with revenue and marketing productivity gains, while fraud detection and risk scoring become central to sustaining digital commerce growth amid rising attack sophistication. However, consumer resistance to automated service, data governance constraints, and “agentic AI” project failure risks create a non-linear ROI landscape that favors selective deployment, strong measurement discipline, and governance-by-design. The Indian ecosystem—characterized by rapid growth, platform competition, and expanding digital public infrastructure—shows accelerated opportunity for AI in vernacular discovery, logistics, trust and safety, and credit enablement, alongside heightened compliance obligations under evolving data protection and consumer protection rules. The paper concludes with a practical research-grounded framework for AI investment prioritization, measurement, and governance in e-commerce.

Keywords: artificial intelligence, e-commerce, personalization, recommender systems, fraud detection, logistics



optimization, India, digital trust, data governance, generative AI

1. Introduction

E-commerce has always been “data-intensive,” but it has not always been “prediction-native.” In earlier phases, advantage in online retail primarily accrued to firms that could aggregate product variety, traffic, and fulfillment capacity. The contemporary competitive frontier is different: advantage increasingly depends on the ability to infer intent, personalize value, allocate inventory dynamically, detect risk in real time, and coordinate last-mile execution under tight time constraints. These tasks are fundamentally probabilistic. AI therefore changes the production function of e-commerce by converting behavioral traces—searches, clicks, dwell time, location, device context, payment signals, returns history—into decisions that can be executed at scale and updated continuously.

The macro context clarifies why AI has become structurally important. UNCTAD has estimated that business e-commerce sales reached **\$27 trillion in 2022** across a large subset of economies, underscoring the magnitude of digital ordering beyond consumer-facing retail alone. At the consumer end, global B2C e-commerce revenues are projected to rise to **\$5.5 trillion by 2027** (U.S. International Trade Administration summary), reflecting continued channel shift and deepening digital penetration. In India, policy-linked digitization, platform competition, and fast adoption in tier-2/3 cities are contributing to a growth curve where the e-commerce market is estimated at **US\$125 billion in**

2024 and projected to reach **US\$345 billion by 2030** (IBEF). As digital commerce scales, operational frictions and trust risks become binding constraints; the average documented cart abandonment rate remains around **70.22%**, showing persistent purchase friction and unrealized demand. At the same time, fraud is rising: Juniper Research projects e-commerce fraud to increase from **\$44.3 billion in 2024** to **\$107 billion in 2029**, a trajectory that directly threatens conversion, payment acceptance, and consumer confidence.

Within this environment, AI is not a single tool but a layered capability stack: predictive models, real-time decision engines, language interfaces, vision-based compliance, and increasingly, generative systems that can produce content and automate knowledge work. The challenge for research—and for journal reviewers—is to move beyond cataloging applications and instead explain **cause–effect relationships**: how AI alters measurable outcomes (conversion, revenue, costs, returns, fraud loss, customer satisfaction), under what conditions, and with what risks and governance requirements.

This paper addresses that gap using an evidence-supported synthesis approach. It focuses on four research questions:

1. **Where** in the e-commerce value chain does AI create the most measurable performance uplift?
2. **How** do adoption-and-scale signals differ between global and Indian contexts, and what structural factors explain divergence?



3. **What** trade-offs emerge between growth optimization and trust/safety (privacy, bias, transparency, fraud, consumer acceptance of automation)?
4. **Which** governance mechanisms (regulatory and organizational) are necessary to make AI-enabled commerce sustainable?

2. Literature-Informed Context and Research Gaps

The research landscape on AI in e-commerce spans multiple traditions: information systems (technology adoption and business value), marketing science (personalization and customer lifetime value), operations research (inventory and routing optimization), and cybersecurity/fintech (fraud detection and authentication). While these literatures offer rich micro-insights, three gaps often appear when practitioners interpret them for e-commerce strategy.

First, value claims are frequently de-contextualized. For example, personalization is widely cited as beneficial, but rigorous interpretations require specifying *what kind* of personalization (recommendations, offers, content, timing), *what data regime* (first-party versus third-party), and *what execution maturity* (measurement, experimentation, and governance). McKinsey's synthesis suggests personalization often drives **5–15% revenue lift** and can increase marketing ROI by **10–30%**, but also highlights a wide

range dependent on sector and execution capability.

Second, trust and adoption constraints are under-modeled. Consumer acceptance of AI-mediated interactions is not guaranteed. Gartner reported that **64% of customers would prefer companies did not use AI for customer service**, and **53%** would consider switching if they learned AI would be used for service, implying that automation can create churn risk when deployed without design safeguards and escalation paths.

Third, macro risk is rising alongside capability. The same generative capabilities used to automate service and content can be weaponized for scams, synthetic identity fraud, and social engineering. Policy and governance therefore become part of the ROI equation. This is increasingly formalized: India's Digital Personal Data Protection Act, 2023 (DPDP) embeds consent, purpose limitation, and withdrawal mechanisms relevant to e-commerce data processing. The Consumer Protection (E-Commerce) Rules, 2020 also impose transparency and consumer-right safeguards affecting platform operations. Internationally, the EU AI Act introduces a risk-tiered compliance regime and staged obligations that will affect cross-border commerce technology and AI vendors.

The implication is that "AI in e-commerce" should be analyzed as a **system**—spanning customer experience, operations, payments and risk, and governance—rather than as isolated applications.



3. Methodology: Secondary-Data Evidence Synthesis

This study uses a structured secondary-data synthesis method designed for research questions that require integration across heterogeneous sources (market size indicators, policy documents, industry research, and operational benchmarks). The approach consists of four steps:

1. **Source selection and reliability weighting.** Priority is given to multilateral and governmental publications (UNCTAD, OECD, MeitY, Department of Consumer Affairs), major consulting/industry research (McKinsey, Accenture, Gartner), and established benchmark organizations (Baymard) along with reputable market intelligence summaries for contextual sizing (e.g., Grand View Research).
2. **Evidence mapping by value-chain function.** Findings are coded into a functional map: discovery → conversion → payment →

fulfillment → returns → retention, plus cross-cutting trust and governance.

3. **Comparative synthesis and normalization.** Where sources use different units, results are normalized into interpretable patterns (e.g., growth rates, indexed adoption signals, and performance ranges). The paper does not present primary survey results; instead it derives insights by triangulating consistent directional signals across sources.
4. **Analytical tables and interpretive narratives.** Tables are used only where they add synthesis value: capability-to-function mapping, adoption and scale indicators, pre/post performance patterns, and risk-mitigation trade-offs. Each table is followed by a detailed interpretation linking the evidence to research arguments.

4. AI Capability Stack and Its E-Commerce Decision Points

Table 1. AI capabilities mapped to e-commerce functions and measurable KPIs

AI capability	Typical e-commerce implementation	Primary mechanism	KPIs most directly affected	Where value concentrates
Machine learning (predictive / ranking)	Search ranking, recommendations, propensity scoring, churn prediction	Converts behavior signals into probabilistic rankings	Conversion rate, AOV, CTR, repeat rate, CLV	High-traffic discovery and basket-building
NLP (incl. multilingual)	Intent detection, query understanding, review	Reduces semantic mismatch; extracts	Search success rate, time-to-product, return rate, CX scores	Catalog-heavy marketplaces; vernacular markets



AI capability	Typical e-commerce implementation	Primary mechanism	KPIs most directly affected	Where value concentrates
	mining, automated merchandising	product/issue themes		
Computer vision	Visual search, image moderation, counterfeit detection, attribute extraction	Converts image data to structured attributes and risk flags	Fraud/counterfeit incidence, listing quality, time-to-list	Trust & safety; fashion/marketplace listings
Optimization + forecasting	Demand forecasting, inventory allocation, route/slot optimization, dynamic pricing	Balances cost/service trade-offs under constraints	Stockouts, fill rate, delivery time, margin	Fast delivery, high SKU volatility categories
Generative AI (LLMs)	Product content generation, conversational commerce, agent assist, policy summarization	Automates language production and knowledge access	Content velocity, service productivity, conversion support	Long-tail SKU catalogs; service-heavy verticals
Anomaly detection / graph ML	Payment fraud, account takeover detection, fake review networks	Detects non-stationary adversarial patterns	Fraud loss rate, chargebacks, false decline rate	Payments, promotions, marketplace integrity

Interpretation of Table 1. The table highlights a critical point: AI creates value where e-commerce decisions are (i) frequent, (ii) uncertain, (iii) sensitive to personalization or risk, and (iv) measurable. Discovery and conversion have long been AI-friendly because incremental improvements scale across millions of sessions. However, the inclusion of anomaly detection and computer vision demonstrates that AI’s role is not only “selling more” but also “preventing loss and protecting trust.” This becomes central when fraud is projected to rise sharply (e.g.,

Juniper’s \$44.3B in 2024 to \$107B in 2029).

The operational emphasis differs by market. In mature markets, recommender systems and customer analytics often dominate the AI agenda. In India, multilingual NLP and trust infrastructure become equally strategic because consumer acquisition increasingly comes from heterogeneous linguistic contexts and varied digital literacy. The mechanism is not cultural; it is informational: when language and catalog structure diverge, the cost of mismatch rises, and NLP becomes a conversion lever. This provides a more



causal explanation than generic claims that “India is diverse.”

5. Market Scale, Adoption Signals, and the India–Global Contrast

Table 2. Macro indicators linking e-commerce scale to AI adoption signals (selected secondary evidence)

Indicator	Global signal	India signal	Why it matters for AI deployment
Total e-commerce scale	UNCTAD estimate: \$27T business e-commerce sales (2022)	High-growth retail e-commerce trajectory; rapid platform expansion	Larger transaction volume increases payoff from small accuracy gains and automation
B2C growth projection	Global B2C e-commerce to \$5.5T by 2027	Market projected \$125B (2024) → \$345B (2030)	Growth increases experimentation budgets and data scale, but also complexity
AI-in-retail market growth	AI in retail estimated \$11.61B (2024) → \$40.74B (2030)	India: accelerated use in fraud detection and platform automation narratives	Indicates investment shift from pilots to infrastructure and productization
Persistent friction baseline	Avg cart abandonment ~70.22%	Similar frictions amplified by COD/returns and logistics variability (industry pattern)	AI gains are partly “friction harvesting”: reducing drop-offs, improving trust
Fraud pressure	Fraud projected \$44.3B (2024) → \$107B (2029)	Payments ecosystem highlights AI as double-edged: efficiency + fraud risk	Risk models become core to sustaining conversion and payment acceptance
Enterprise AI maturity signal	AI-led processes: 9% (2023) → 16% (2024) (Accenture)	Rapid adoption pressure in consumer internet sectors and quick commerce	Suggests “capability gap” between leaders and laggards; ROI depends on maturity

Interpretation of Table 2. Table 2 connects scale and adoption—not as abstract “digital transformation,” but as an economic argument. When e-commerce is measured in trillions of dollars, even small improvements in ranking relevance, fraud detection precision, or fulfillment efficiency create material value. This is

why AI-in-retail investment is projected to grow quickly in global estimates.

The India signal is not only market size; it is *market structure*. India’s e-commerce trajectory (IBEF) implies rapid demand growth, but also higher variance in delivery constraints, payment preferences, and language contexts. This increases the



marginal value of AI systems that can adapt locally (vernacular NLP, dynamic ETA prediction, fraud and returns scoring). Meanwhile, the persistence of cart abandonment near 70% globally demonstrates that a large portion of demand remains unrealized due to friction. AI becomes a “friction-reduction technology” as much as a “growth technology,” which is a more defensible claim because it uses a measurable baseline.

Finally, Table 2 surfaces an often-missed inference: AI investment payoffs will likely be **uneven** because enterprise AI maturity is uneven. Accenture’s data (AI-led processes 9%→16% in one year) implies a widening performance gap between organizations that can redesign processes and those that merely add AI tools. For e-commerce, this distinction matters because model outputs only create value when embedded into decisions (pricing, promotions, fraud blocks, inventory moves) with governance and measurement.

6. Performance Effects Across the E-Commerce Funnel: What Changes and Why

6.1 Discovery and Conversion: Personalization and Recommendation as Revenue Mechanisms

Personalization is often described as “showing the right product to the right customer,” but analytically it is better framed as reducing entropy in demand. Consumers arrive with partial intent; search and recommendation systems convert that partial intent into a product set with higher expected utility. The expected business

outcome is not merely higher click-through; it is reduced search cost, higher basket completion, and improved long-term retention because customers learn that the platform “understands” them.

A key evidence anchor is McKinsey’s synthesis that personalization can reduce customer acquisition costs substantially and is associated with **5–15% revenue lift** and **10–30% marketing ROI improvement**, with variation across execution contexts. These ranges are particularly relevant because they come from multi-industry observations rather than single-case marketing claims. They also imply a methodological caution: personalization is not deterministic; it is conditional on data quality, experimentation, and channel integration.

At the same time, platforms and commerce enablement firms increasingly claim large uplifts from AI recommendations. For instance, Shopify’s AI statistics page (a synthesized industry-facing summary) asserts that smart recommendations can more than double conversion and lift order value. For a reviewer, such claims should be treated as directional rather than universal; yet they align with the mechanism that recommendation engines improve relevance and cross-sell effectiveness. The academically defensible stance is to interpret these as **upper-bound outcomes** for well-instrumented merchants and compare them against more conservative consulting ranges.



6.2 Checkout and Friction: AI as a “Drop-off Minimization” System

Cart abandonment is one of the most stubborn empirical realities of e-commerce. Baymard’s aggregated benchmark suggests an average abandonment rate around **70.22%**, implying that most sessions that express purchase intent do not convert. This fact changes how AI value should be modeled: if the funnel is leaky, AI ROI can come from *stopping leakage* as much as from acquiring new users.

AI-driven interventions include: real-time risk-based authentication (reducing false declines), personalized shipping/return messaging, predictive delivery windows, dynamic bundling, and fraud-aware promotion eligibility. The causal channel is reduction of perceived risk and effort at the moment of payment and shipping commitment.

6.3 Trust, Fraud, and Payments: AI as the Invisible Infrastructure of Conversion

Fraud is not just a cost line; it is a conversion determinant. Aggressive fraud controls can increase false declines, which reduces legitimate orders and damages consumer trust. Weak fraud controls increase chargebacks, merchant losses, and platform integrity problems. Juniper’s projected rise in e-commerce fraud value from **\$44.3B (2024) to \$107B (2029)** provides a macro signal that adversarial pressure is intensifying.

In India, the payments ecosystem is explicitly framed as a double-edged technology environment: innovation increases efficiency but also expands the

attack surface. PwC’s Indian payments handbook notes that generative AI can improve speed and experience while also “providing fuel” to payment fraud, reinforcing the need for advanced fraud detection and governance. This aligns with the operational reality that fraud detection must become adaptive (anomaly detection, graph models, device fingerprinting, behavioral biometrics) and must be continuously retrained and monitored.

6.4 Customer Service: Productivity Gains Versus Consumer Resistance

Customer service is one of the most tempting areas for AI automation due to direct labor costs and high ticket volume. Yet the acceptance constraint is material. Gartner’s finding that **64% of customers would prefer companies didn’t use AI for customer service** introduces a measurable demand-side risk. This suggests that AI service deployments that prioritize deflection and cost reduction without maintaining human escalation paths may increase churn.

The more robust value proposition is **agent augmentation** rather than full replacement: AI drafts responses, summarizes context, retrieves policy, and guides human agents. The productivity upside is consistent with broader enterprise operations research showing AI-led processes correlate with higher productivity and revenue growth (Accenture). However, the governance requirement becomes stronger: service AI must be measured on resolution quality, not merely deflection rate.



7. Pre-AI vs AI-Enabled Outcome Patterns

Table 3. Performance pattern comparisons (pre-AI vs AI-enabled) using triangulated secondary evidence

E-commerce area	Pre-AI baseline pattern	AI-enabled pattern (evidence-supported ranges/signals)	Key evidence anchors
Personalization & marketing efficiency	Broad segmentation; higher acquisition costs; less precise targeting	Revenue lift often 5–15% ; marketing ROI +10–30% (execution-dependent)	McKinsey synthesis
Funnel friction (checkout)	Persistent abandonment ~70%; limited real-time adaptation	AI targets friction points (risk-based flows, personalized reassurance, better relevance); value is “leakage reduction”	Baymard abandonment benchmark
Fraud & trust	Rule-based controls; brittle to new attack patterns; costly false positives	Adaptive models + anomaly detection; necessity rises as fraud projected \$44.3B→\$107B (2024–2029)	Juniper fraud projection
Customer service	High ticket volume; variable quality; limited 24/7 coverage	Automation/augmentation improves productivity but faces resistance: 64% prefer no AI in service	Gartner consumer preference
Enterprise operational maturity	Automation uneven; manual processes dominate in many firms	AI-led processes increased 9%→16% (2023–2024) ; leaders outperform on growth and productivity	Accenture operations research

Interpretation of Table 3. The table makes a core analytical contribution: it separates AI’s impact into (1) **growth and efficiency** effects and (2) **trust and adoption** constraints. Personalization gains are supported by a stable consulting synthesis (McKinsey), but the magnitude is conditional—suggesting that reviewers should interpret uplift claims as *capability-dependent* rather than universal.

The checkout friction baseline is crucial. A 70% abandonment rate is not a marginal optimization problem; it is a structural gap between intent and completion. AI therefore creates value by identifying micro-frictions and tailoring interventions, but it also raises measurement standards: firms must attribute uplift correctly to avoid confusing correlation (e.g., high-intent users respond more to interventions) with causation.



Fraud and trust are shown as a *growing* value zone because the threat environment is worsening. This means AI spending in e-commerce cannot be justified only by incremental sales; it must be justified by *preserving* sales that would otherwise be lost to fraud controls, chargebacks, and consumer fear. Finally, customer service is presented as the domain where AI enthusiasm is most likely to overrun consumer preferences, creating a realistic limitation that a credible research paper must acknowledge.

8. The Indian E-Commerce Ecosystem: AI Levers and Structural Constraints

India's e-commerce growth curve is frequently described with market size projections, but the research-relevant question is **which structural features make AI especially valuable or risky in India.**

8.1 Growth, heterogeneity, and the vernacular frontier.

IBEF's projection (US\$125B in 2024 to US\$345B in 2030) implies rapid customer acquisition beyond metro elites. As adoption diffuses, language and digital literacy heterogeneity becomes more important. AI can reduce friction through multilingual query understanding, voice-assisted commerce, and improved product attribute normalization. The causal channel is reduced cognitive effort and improved matching between intent and catalog.

8.2 Logistics and quick commerce: prediction as an operational necessity.

Quick commerce growth in India is intensifying expectations around delivery speed and reliability, creating a premium on accurate demand forecasting and micro-fulfillment optimization. Reuters reported that quick commerce contributed to a major share of e-grocery orders and expanded rapidly, illustrating how delivery-time competition compresses operational slack. AI-driven routing, slotting, and inventory allocation become not just cost optimizers but brand differentiators because "speed" is experienced as trust.

8.3 Trust, fraud, and the payments interface.

India's payments digitization increases the need for robust fraud controls. The payments literature emphasizes the dual-use nature of advanced AI: while it improves processing efficiency, it can also amplify fraud techniques. E-commerce platforms therefore need risk models that are adaptive, privacy-aware, and integrated into checkout flows without harming legitimate users.

8.4 Governance and compliance are not optional.

India's DPDP Act creates explicit requirements around consent, withdrawal, and lawful processing that affect recommendation engines, targeted advertising, and profiling. Additionally, the Consumer Protection (E-Commerce) Rules, 2020 impose transparency and consumer safeguards relevant to product information, refunds, and grievance mechanisms. For AI systems, the compliance task is not merely "privacy policy updates"; it is architectural: data minimization, purpose limitation, and



auditable decision logic become necessary design constraints.

9. Risks, Limitations, and Governance-by-Design

Table 4. AI risks in e-commerce and evidence-grounded mitigation approaches

Risk category	How it manifests in e-commerce	Why it is increasing	Mitigation direction (governance + technical)	Evidence anchors
Consumer trust backlash	Users dislike AI-only service; churn risk	Consumers prefer human options; fear of job loss or bad resolution	Hybrid service design, human escalation, transparency, quality metrics beyond deflection	Gartner: 64% prefer no AI for service
Fraud arms race	Synthetic identity, ATO, promotion abuse, deepfake-assisted scams	Fraud projected to rise sharply; gen-AI lowers attacker cost	Adaptive anomaly detection, graph ML, stepped-up authentication, continuous monitoring	Juniper fraud growth ; payments risk framing
Privacy and lawful processing constraints	Over-collection, unclear consent, profiling without safeguards	Regulation and public scrutiny rising	DPDP-compliant consent flows, purpose limitation, data minimization, privacy-preserving ML where needed	DPDP Act text
Regulatory fragmentation	Cross-border AI compliance complexity	EU AI Act staged obligations; global governance divergence	Risk classification, documentation, vendor due diligence, model cards and audits	EU AI Act framework
Hype-driven deployment failure	High spend with unclear outcomes; “agent washing”; pilot purgatory	Rapid vendor proliferation; unclear ROI measures	Use-case prioritization, strict experimentation, post-deployment monitoring, kill-switches	Gartner/Reuters on agentic AI project cancellations
Experience degradation via over-automation	Poor answers, looping bots, unresolved issues	Pressure to cut costs; measurement mis-specified	Evaluate on resolution, CSAT, complaint rates; maintain human-first for complex cases	Gartner preference + service resistance



Interpretation of Table 4. Table 4 emphasizes that AI risk is not an “add-on” section; it is endogenous to e-commerce value creation. The same automation that reduces cost can degrade experience, and consumer preference data indicates that over-automation can actively destroy value through switching behavior. Similarly, fraud is not only a security function; it is a growth constraint, because stricter controls can reduce authorization rates and increase abandonment—especially when consumers already face high friction at checkout. The projected fraud escalation (Juniper) makes it likely that platforms will increase controls, raising the importance of AI models that can reduce false positives while catching new threats.

Regulatory constraints should be treated as design requirements. DPDP introduces consent withdrawal mechanics and obligations that influence how platforms train models, store profiles, and deliver targeted offers. Internationally, the EU AI Act’s risk-based regime and staged obligations affect e-commerce firms indirectly via vendor compliance and cross-border data flows.

Finally, the “hype risk” is empirically relevant. Reuters’ coverage of Gartner’s view that a large share of agentic AI projects may be scrapped reflects a broader reality: e-commerce AI must be justified with measurable decision impact, not “AI presence.” This aligns with the maturity gap observed in enterprise AI-led processes: only a minority of firms have fully modernized workflows, implying that many deployments will underperform due

to organizational, not algorithmic, limitations.

10. Discussion: A Causal Model of How AI Changes E-Commerce Outcomes

Integrating the evidence, AI’s impact in e-commerce can be described through five causal pathways:

1. **Relevance pathway (demand matching):** ML and NLP increase match quality between intent and catalog → higher conversion and AOV → higher CLV. Supported by personalization uplift ranges.
2. **Friction pathway (effort and uncertainty reduction):** AI reduces checkout and decision friction through better information, risk-based flows, and predictive delivery confidence → lower abandonment and higher completion. Baseline urgency indicated by persistent abandonment.
3. **Trust pathway (risk control without harming good users):** AI improves fraud detection precision and speed → protects margins and reduces consumer fear; but only if false positives are controlled. Threat growth indicated by fraud projections.
4. **Productivity pathway (automation and augmentation):** AI shifts human labor from repetitive tasks to exception handling and creative work →



productivity gains; but consumer acceptance constraints limit full automation in service.

- Governance pathway (compliance and legitimacy):** Data protection and consumer rules constrain data use and automated decisioning → AI must be designed with transparency, consent, and auditability → sustainable deployment.

These pathways explain why AI ROI is highly **context-dependent**. A high-traffic marketplace with weak catalog structure may see more benefit from NLP-based attribute extraction than from generative content. A quick commerce player may find that demand forecasting and routing optimization produce more defensible advantage than a chatbot. A regulated environment may reduce the feasible feature set for personalization unless consent and purpose limitation are robustly engineered.

11. Implications for Research and Practice

For researchers, the evidence suggests three priority directions:

- Attribution discipline:** Many AI uplift claims confound selection effects. Future studies should emphasize experimentation designs (A/B tests, multi-armed bandits with guardrails) and long-horizon outcomes (repeat rate, churn, complaint rates), not only short-term conversion.

- Trust–growth co-modeling:** Fraud controls, privacy constraints, and consumer acceptance should be modeled jointly with revenue outcomes, because they interact with conversion and retention.
- India-specific mechanisms:** Research should focus on vernacular NLP, COD/returns behavior, and trust infrastructure interactions—areas where generic Western e-commerce models may not transfer.

For practitioners, the practical lesson is that AI strategy should be organized around measurable decision points: ranking, pricing, fraud acceptance, inventory moves, and service resolution. The highest-performing organizations appear to be those that redesign processes rather than layer AI on top of old workflows—consistent with evidence that AI-led processes correlate with superior outcomes.

12. Conclusion

AI has become a structural capability in e-commerce because the core problems of digital commerce—matching, friction, trust, and coordination—are prediction and decision problems at scale. Secondary evidence indicates that personalization is often associated with meaningful revenue and marketing ROI uplift, while persistent cart abandonment underscores large unrealized gains from friction reduction. At the same time, rapidly rising fraud projections imply that AI-based risk control is not optional but foundational to protecting both margins and consumer confidence.



The Indian e-commerce ecosystem amplifies AI's relevance due to high growth, linguistic heterogeneity, and operational complexity, while also increasing the importance of governance under the DPDP Act and consumer protection rules. Internationally, regulatory fragmentation (e.g., EU AI Act) and the empirical risk of hype-driven project failure reinforce that sustainable AI in e-commerce requires measurement discipline, governance-by-design, and a selective deployment strategy focused on high-leverage decisions.

In sum, AI is best understood not as a set of features, but as an operating system for digital commerce—one that can increase value creation only when aligned with trust, compliance, and human-centered experience design.

References

1. Accenture. (2024). *New Accenture research finds that companies with AI-led processes outperform peers* (Press release).
2. Baymard Institute. (2025). *Cart abandonment rate statistics (average documented rate 70.22%)*.
3. European Commission. (2026). *AI Act: Shaping Europe's digital future* (policy page).
4. Gartner. (2024). *Gartner survey finds 64% of customers would prefer companies didn't use AI for customer service* (Press release).
5. India Brand Equity Foundation (IBEF). (2025). *Indian e-commerce industry analysis / presentation*.
6. Juniper Research. (2024). *eCommerce fraud to exceed \$107 billion in 2029* (Press release).
7. McKinsey & Company. (2021). *The value of getting personalization right—or wrong—is multiplying*.
8. McKinsey & Company. (2023). *What is personalization?* (Explainer/PDF).
9. MeitY, Government of India. (2023/2024). *The Digital Personal Data Protection Act, 2023* (Act text).
10. OECD. (2019). *OECD AI Principles / Recommendation of the Council on Artificial Intelligence*.
11. PwC India. (2024). *The Indian payments handbook 2024–2029* (Report PDF).
12. Reuters. (2025). *Over 40% of agentic AI projects will be scrapped by 2027, Gartner says* (News report).
13. UNCTAD. (2023). *Measuring the value of e-commerce* (Report).
14. UNCTAD (via UNESCAP/UNCTAD presentation referencing UNCTAD estimate). (2022).
15. *Digital trade and e-commerce for development* (slide).
16. U.S. International Trade Administration. (2024). *eCommerce sales & size forecast* (summary page).