SWOT Assessment of Agentic Commerce for Retailers

edgar, dunn & COMPANY

borderless payment consulting

January 2026

Executive Summary

The move to agentic commerce represents the most significant shift in retail since the start of e-commerce. Success requires moving beyond "chatbots" to a system where your brand can communicate, negotiate, and transact directly with consumer-led AI agents. By addressing the 24 critical considerations identified across this SWOT assessment, retailers can transition from passive storefronts to active participants in an agent-led economy.

To master these 24 considerations, retailers should prioritise interoperability and trust. The transition from a web-based "pull" model to an agent-led "push" model is not merely a technical upgrade; it is a total redesign of the customer journey. While this report outlines numerous calls to action across technology, operations, merchandising, and risk functions, retailers must prioritise ruthlessly to avoid spreading resources too thin. Developing a clear, phased roadmap, starting with immediate actions like API optimisation and product data structuring, followed by mid-term investments in auditable logs and “Know your Agent” (KYA) compliance, is essential for building competitive advantage without attempting to achieve everything simultaneously. This structured approach ensures focus on high-impact capabilities first, delivering measurable progress while maintaining flexibility to adapt to rapidly evolving agentic commerce standards and adoption curves.

Agentic commerce represents a substantial disruption within digital retail, with the global opportunity estimated at approximately USD 2.9 billion by 2030. Unlike the gradual shift from physical stores to e-commerce in the late 1990s and early 2000s, agent-led shopping is expected to cannibalise traditional e-commerce at a materially faster rate, as AI agents increasingly intermediate product discovery, comparison, and checkout. This dynamic is fundamentally different from omnichannel, which blurred the boundaries between offline and online sales; agentic commerce instead reconfigures how online demand itself is generated and directed, reshaping the economics and competitive landscape of e-commerce.

Retailers who successfully design for Agent-to-Agent commerce will capture the "intent" of the consumer long before they ever reach a traditional search engine or social media feed. Agentic commerce is an exciting opportunity for retailers.

Report Contributors

Mark Beresford
Director

Davide Villa
Manager

Elisabetta Nadal
Consultant

Reuben Joseph
Business Analyst

Table of Contents

Click the link to go to the section

edgar, dunn & COMPANY

Introduction

The purpose of a SWOT assessment of agentic commerce is to provide a concise strategic view of its viability and impact by clarifying the strengths of autonomous, AI-driven transactions, exposing internal weaknesses such as reliability, data readiness, and governance gaps. Identifying opportunities for new business models and market expansion through Agent-to-Agent interactions and predictive purchasing, and exploring external threats including regulatory uncertainty, security risks, fraud, platform dependence, and loyalty erosion.

Overall, this SWOT Assessment aims to help executives in retail to make informed investments, define product roadmaps, and create risk-mitigation plans by balancing value creation against operational, legal, and reputational considerations.

While a SWOT assessment is a staple of strategic planning, its simplicity is both its greatest strength and its primary limitation. It is often criticised for being a "snapshot in time". Agentic commerce is evolving at a remarkably fast pace, moving from early experimentation in 2024 to real-world deployments and measurable traffic growth in 2025 and early 2026. The adoption is being driven by major payment platforms, fintech companies, embedding autonomous purchasing capabilities (such as chat-based checkouts and AI shopping agents), with retail traffic from AI interfaces growing exponentially year-over-year, if not month-over-month, and early use cases already visible on major retail commerce sites. Edgar, Dunn & Company (EDC) forecasts suggest the market could expand into the billions of dollars by 2030, with projections of significant global sales and widespread consumer use, highlighting that what was once conceptual innovation is quickly becoming a mainstream business reality.

The reader must recognise this assessment is a snapshot in time. The threats and weaknesses of agentic commerce today will change in the next 6 to 12-months. Just as we have seen rapid adoption, the strengths and opportunities of agentic commerce in retailing are likely to surprise us all. As we have seen the mass adoption of contactless payment at the Point of Sale took around eight years. The adoption of ChatGPT has scaled from a limited proof of concept into one of the world's most used online tools in just a few years. At the end of 2025, it had around 800 million weekly active users. ChatGPT is not the only AI tool that is expected to stimulate and support growth in agentic commerce. The early adoption and growth of agentic commerce is expected to be equally significant. Retailers and payment service providers must put in place a plan for agentic commerce. This report aims to be your initial point of reference.

Notes on data and information sources

This report draws on a rigorous, multi-method research approach combining primary qualitative insights from in-depth interviews with enterprise retail clients and secondary desk research, ensuring a balanced, comprehensive perspective on the trends in agentic commerce.

Interviews with senior stakeholders from major retailers and fintech providers provided firsthand accounts of implementation challenges, adoption barriers, and early ROI signals, capturing nuanced operational realities not yet reflected in publicly available data. Desk research was systematically expanded to include analysis of recent agentic commerce announcements by leading retailers (e.g., pilots for autonomous checkout and inventory agents) and fintech companies (e.g., payment protocols like Mastercard Agent Pay, Shopify and Stripe's agent APIs), alongside benchmarking against industry forecasts, industry updates, and agentic announcements to triangulate findings and enhance this report's reliability.

The methodology in preparing this SWOT assessment provides actionable, evidence-based sizing and a comprehensive list of calls to action for retailers that is grounded in both practitioner experience, Edgar, Dunn & Company (EDC) expertise and forward-looking market developments.

Table of Contents

Click the link to go to the section

edgar, dunn & COMPANY

Overview of Agentic Commerce in the retail sector

Agentic commerce in retail is the use of AI "shopping agents" that autonomously search, compare, decide, and complete purchases (including post-purchase tasks) on behalf of customers across retailers and channels, within parameters the customer sets. These agents effectively become the primary interface to the retail ecosystem, rather than the retailer's own website or app.

Agentic commerce will shift retail customer journeys from manual, channel-by-channel shopping to intent-driven, delegated journeys where AI agents handle most of the work. Customers will increasingly brief an AI agent ("I need a new shirt for a wedding" or "top up household essentials – for the kitchen and bathrooms") rather than browsing websites and apps themselves. Agents will interpret signals (such as calendar events, past purchases, preferences, budgets, the weather, time of the year) to transform intent into a shopping mission, then run discovery and comparison across retailers. Agents will anticipate customer needs (e.g., noticing low pantry stock or upcoming weather changes) and trigger replenishment or recommendations before the customer explicitly shops.

The discovery, evaluation, and checkout blur into one flow - the AI agent narrows options, checks reviews and prices, and executes payment and shipping choices in the background. The agent will select the best payment method, apply the appropriate loyalty points and coupons, and choose delivery options automatically, massively reducing friction and form-filling.

Shopping journeys become hyper-personalized, reflecting style, fit, dietary needs, brand preferences, and even schedule or mood, learned over time and reused across retailers. Agents will have a memory of your previous purchases and likes and dislikes. The same agent will be able to carry context from store to store and channel to channel (web, app, store, messaging), making the journey feel continuous instead of fragmented by each touchpoint. The winning journeys in the future will depend on rich structured data, clear service levels, predictable returns, and agent-readable loyalty/benefits so that agents choose a retailer based on less about how a site looks to humans and more about how "legible" a retailer is to AI agents.

For a retailer, the question is less "will it happen?" and more "how much transaction volume will shift, and in which categories, by when?". Agentic commerce will happen, it will be gradual, uneven, and initially for journeys that are repetitive and price-sensitive. The direction of travel is clear enough that retail roadmaps now assume agentic use cases rather than treating them as an experiment. Today, there are three leading models that are shaping this new ecosystem in retail as described here:

In Agentic Commerce there are three leading models that are shaping this new ecosystem in retail and many other sectors

Consumer Agents (B2A)

This is where retail businesses interact directly with AI agents representing consumers, tailoring their offerings to algorithmic decision-making. APIs that let the consumer's agent query availability, product prices, sizes, and specifications, delivery arrangements, etc. In this model, the agent would be working for and operated by the consumer. In other words, the entry point for the shopper is their AI agent, such as ChatGPT or Perplexity. Commonly referred to as Business-to-Agent (B2A).

Merchant Agents (A2C)

A2C is where the autonomous AI agents serve or sell directly to shoppers, providing personalised products, services, or recommendations. For example, an Amazon or Walmart has its own AI agent that curates and offers products packages tailored to the customers preferences. In this model, the AI agent will operate under the brand of the retailer. In other words, the entry point for the consumer in this model is not their AI agent but the website or app of the merchant. Commonly referred to as Ageny to Consumer (A2C).

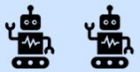

Agent to Agent (A2A)

A2A is where the consumer's agent and the merchant's agent independently negotiate, collaborate, or transact with one another, forming a fully automated layer of commerce (e.g., where a buyer's AI agent communicates with a retailer's agent to autonomously discover products, compare prices, negotiate terms, and execute purchases). This is expected to be normal practice for a weekly grocery shop which could be handled entirely agent-to-agent dialogue. A2A (agent-to-agent) commerce is widely seen as the end goal of the agentic commerce evolution, where autonomous AI agents negotiate, transact, and optimise entirely without human intervention.

Table of Contents

Click the link to go to the section

edgar, dunn & COMPANY

Headline agentic commerce market sizing estimates for retail

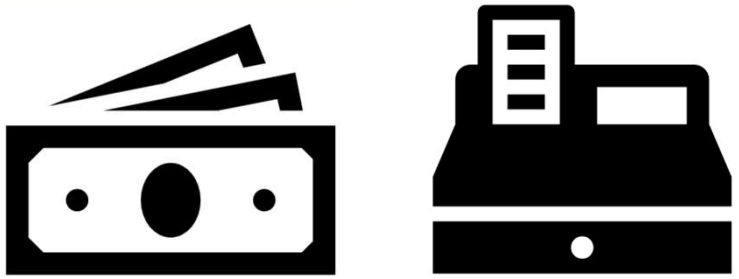

$2.9 billion

Total value of retail sales conducted via AI agents by 2030

29%

Proportion of global retail e-commerce will be completed via AI agents by 2030

Retailers in North America and Asia Pacific will be early adopters

Sizing the Agentic Commerce Market for retail

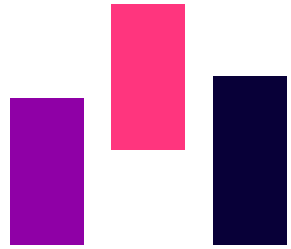

Introduction

The main challenge in sizing the agentic commerce market in retail is that it is an emerging, partially overlapping rather than a clearly segmented category. Historical data and established definitions do not yet exist. This makes it hard to avoid double-counting, such as mixing "agent-influenced" with "agent-executed" retail sales and to separate genuine new demand from a channel shift within existing e-commerce.

Agentic commerce is widely acknowledged as a subset of e-commerce, representing autonomous AI-driven transactions (discovery, negotiation, purchase) within the broader online retail ecosystem rather than a standalone category. This positioning clarifies market sizing by framing it as a channel shift, projected to capture 15-25% of e-commerce volume by 2030, thus avoiding any inflated total addressable market (TAM) estimates.

Methodology

Within this report EDC has cross-checked top-down estimates (share of online retail flowing through agents) with bottom-up use-case modelling (e.g., general retail, grocery, electronics, specialty, etc.) to test whether the numbers are plausible and to avoid over-inflation.

The consumer population has been split into the following 5 generations: Gen Alpha, Gen Z, Millenial, Gen X, and Baby Boomer. Each generation has differing AI usage, based on age and openness to AI[1]. Data has been extracted from World Bank to forecast the proportion of each generation for the 5 major regions (North America, Europe, Asia-Pacific, Latin America and Caribbean, and Middle East and Africa). Gen Alpha proportions start from 2026, to account for the proportion of Baby Boomers gradually decreasing as they age. From this, a total addressable e-commerce transaction volume per consumer generation has been calculated.

The proportion of e-commerce transaction volume via agentic commerce for each consumer generation is expected to differ greatly between generations, and this percentage has been applied to the addressable e-commerce transaction volume to derive the total e-commerce spend for each generation on a global basis.

Each region has been assigned a different "first year" of agentic commerce, based on how ready the economy is for AI agents to carry out e-commerce transactions. Globally, 2025 was the first year of agentic commerce, however it was an emerging technology, and so the proportion of e-commerce transaction volume was estimated at less than 5% for all generations. By 2030, these proportions are expected to increase massively, driven by the high adoption from Gen Alpha and Gen Z.

Spending weights for each generation were applied to get the 2025 market size, reflecting the amount each generation spends through agentic commerce relative to the average consumer.

Note 1: More information on generations and AI usage can be found on the following slide

Sizing the Agentic Commerce Market for retail – Consumer Generations

Consumer Generation[1]	Description
Gen Alpha (1-15 years old)	A consumer who is agentic-native, and who has grown up into e-commerce that is primarily driven by agentic commerce flows.
Gen Z (13-28 years old)	A consumer who enjoys emerging technologies and uses Agentic Commerce solutions with relatively high frequency across all purchase types.
Millennial (29-44 years old)	A consumer who is pragmatic and purposeful, leveraging Agentic Commerce solutions to automate the purchase of routine, low-value transactions and streamline cumbersome recurring purchases.
Gen X (45-60 years old)	A consumer who is research-oriented in nature, utilising Agentic Commerce solutions to facilitate in-depth research prior to purchasing one-off, high-ticket/consideration items.
Baby Boomer (61-80 years old)	A consumer who is either not inherently opposed to emerging technologies and decides to try Agentic Commerce, though does not greatly adopt the solution, or who rejects emerging technologies and does not use Agentic Commerce at all.

edgar, dunn & COMPANY

Note 1: Age is assumed as of 2026
Sources: Bain, Parents

Agentic Commerce as a Proportion of Total e-commerce, Global, 2026 & 2030

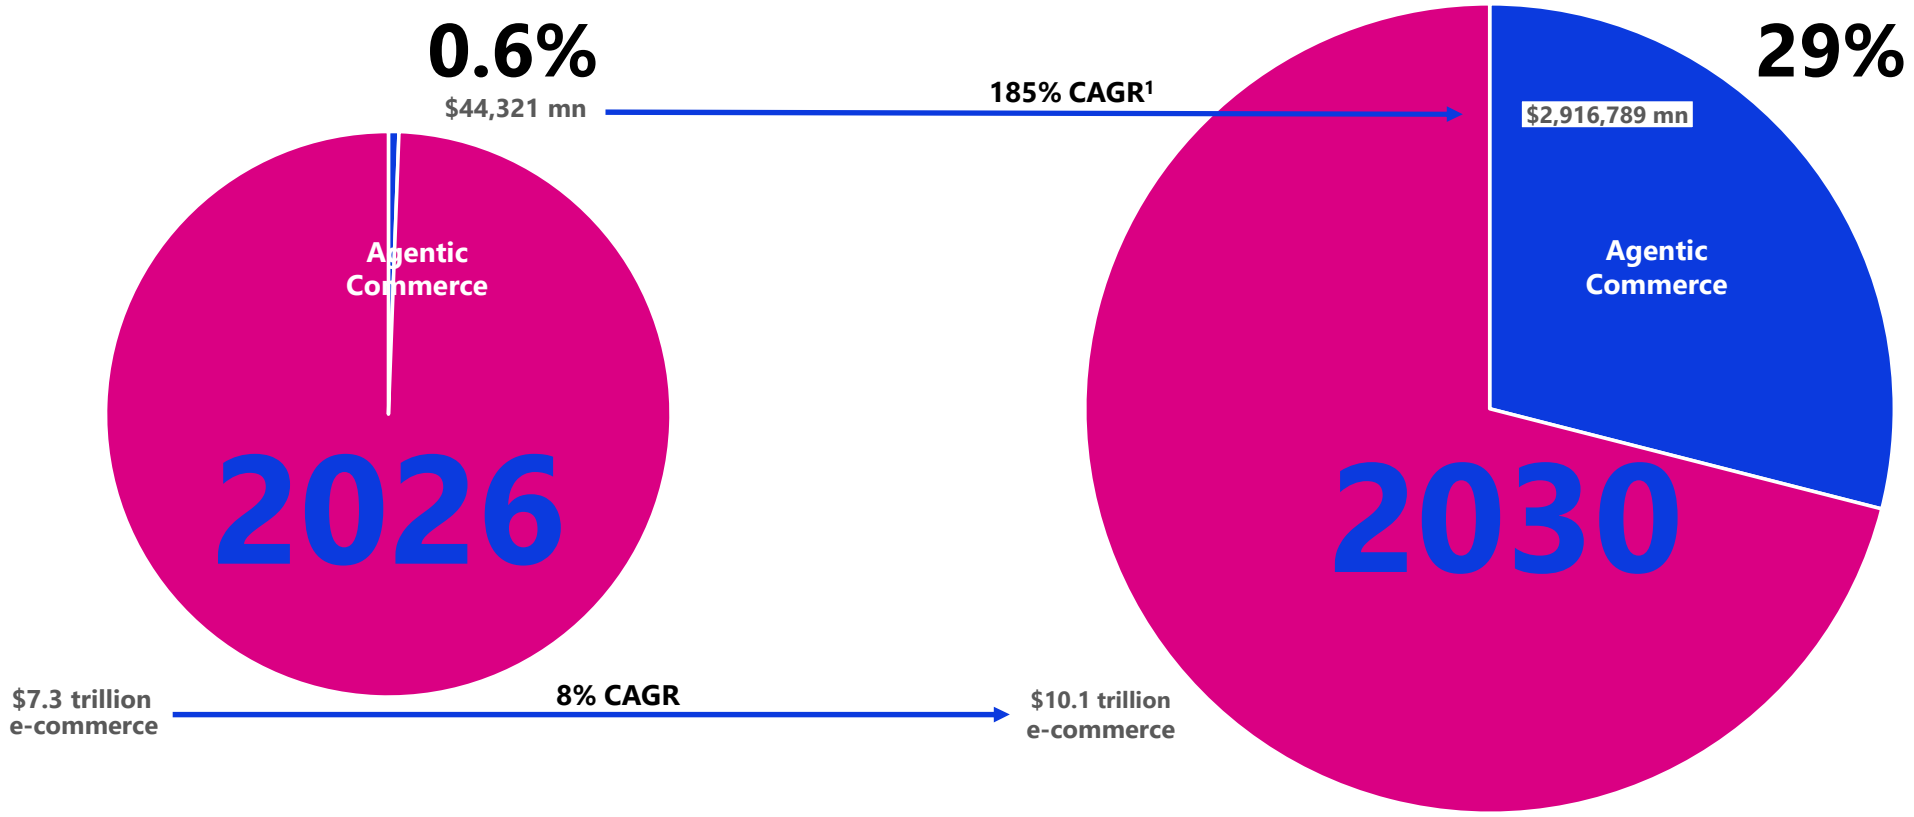

Note 1: 185% CAGR due to very low agentic commerce market share in 2026. CAGRs taken from the following years (2027, 2028, 2029), plateau at 13%, reflecting the large growth expected in the first full year of agentic commerce
Sources: EDC Analysis, World Bank

Agentic Commerce as a Proportion of Total e-commerce, by Region, 2030

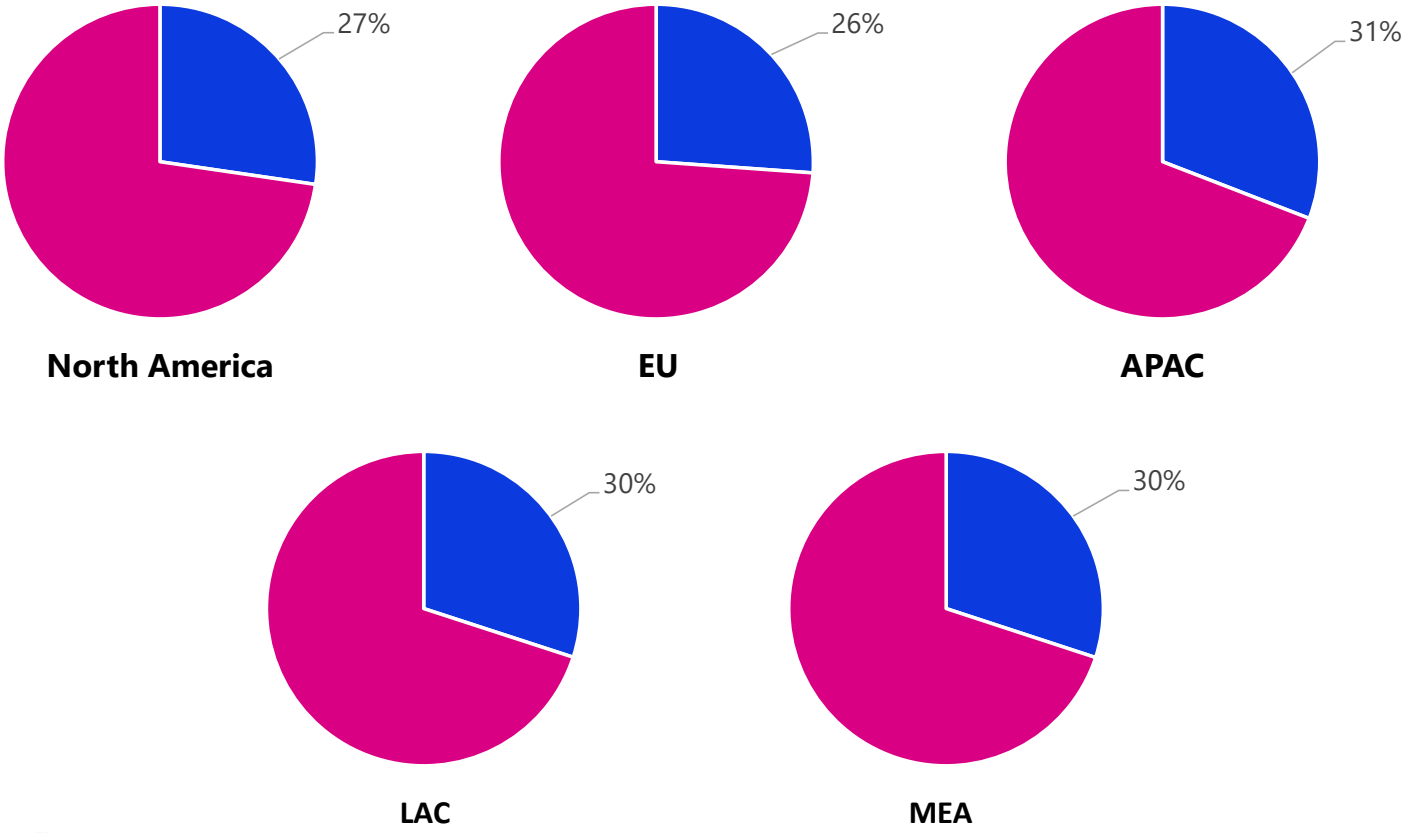

- High retail sales growth in APAC are expected to contribute to the greater adoption of agentic commerce
  - Consumers in APAC tend to be very open technological advancements in payments, as highlighted by the large adoption of super apps like WeChat and Alipay. Therefore, agentic is expected to be adopted by consumers
- Infrastructure variations between the more developed GCC region, and the less developed, larger Africa region result in a lower agentic commerce share for the MEA region by 2030
- Slower AI penetration and a lack of infrastructure compared to other regions, creates barriers to agentic commerce adoption in South & Central America
- Despite differing "first years" of adoption, global market share is at least 26% in all regions by 2030, highlighting the expected growth of agentic commerce
- Adoption gaps will persist, but are expected to narrow over time as the less-developed markets advance their agentic infrastructure

Sources: EDC Analysis, Statista, World Bank

Propensity for AI adoption in retail by subsector

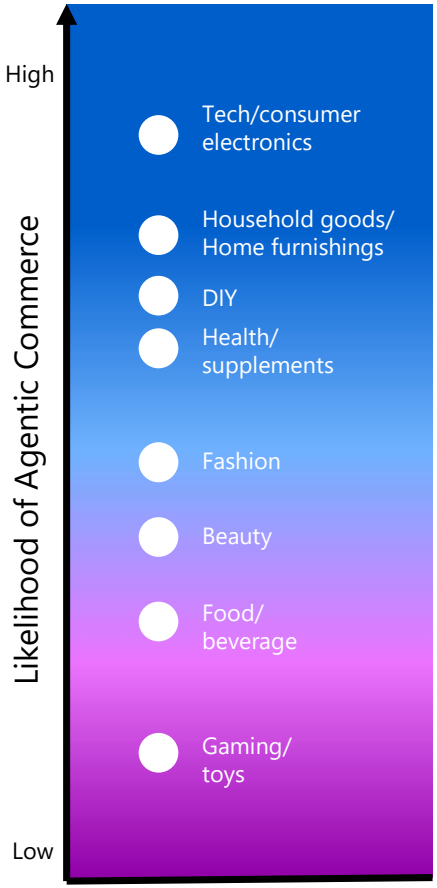

- Consumers show the highest interest in using AI for shopping technology and consumer electronics at 54%, because these products involve complex specifications, rapid innovations, and personalised recommendations that agents excel at simplifying
- Lower adoption in food, beverages, and toys reflects that these purchases are habitual and low-risk, reducing the need for AI agents
  - Consumers already feel confident making these decisions quickly, limiting the benefit of AI agents. Where AI does add value in these categories, it is more likely to be through reminders and bundling, rather than discovery or comparison
- Agentic commerce will fragment retail: scaling first in categories with high decision friction (too many options, too much information, expensive) – not evenly across retail. AI agents are more likely to act as autonomous buyers in tech and travel, and as discovery assistants in everyday retail
- Retailers in high-adoption categories must optimise for agent consumption, by exposing structured data, transparent pricing, and frictionless checkout. In lower-adoption categories, competitive advantage will come from embedding agents into loyalty and replenishment rather than one-off transactions.

Consumers least likely to use AI agents in this vertical

Consumers most likely to use AI agents in this vertical

Sources: EDC Analysis, PartnerCentric

Table of Contents

Click the link to go to the section

edgar, dunn & COMPANY

SWOT Assessment of Agentic Commerce in Retailing

Click to return to table of contents

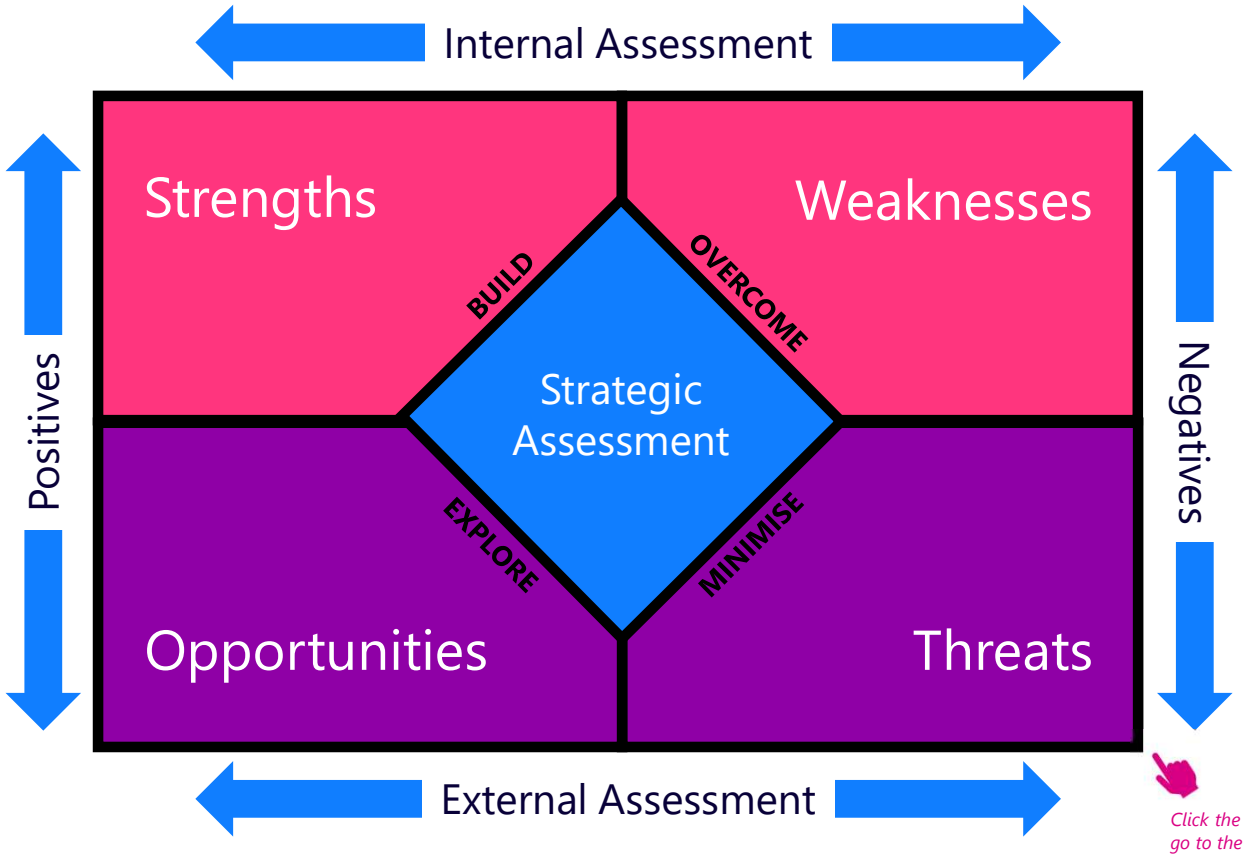

Click the link to go to the section

A SWOT assessment is easy to understand because it uses a simple four-box structure (Strengths, Weaknesses, Opportunities, Threats) that mirrors how people naturally think about what's going well, what isn't, and what could happen.

Summary of the SWOT Assessment of agentic commerce in retailing

Click to return to table of contents

Strengths

- Personalised Engagement
- Operational Efficiency
- Cost Control
- Optimise Checkout & Payment
- Consumer Adoption

Weaknesses

- Reduced Brand Control / Brand Dilution
- Discoverability Risk
- Trust and Transparency
- Mixed baskets
- Multiple Standards and Agentic Protocols
- Unproven commercial models

Opportunities

- Higher Conversion Potential & lower abandonment
- Always-on sales channel
- Smarter demand and inventory management
- New revenue models
- Intelligent upselling & cross-selling
- Agent to Agent efficiency

Threats

- Fraudsters using new attack channels
- Diminished direct customer relationships
- Reliance on third party platforms
- Liability and Dispute Ambiguity
- Margin erosion
- Lack of visibility of customer data
- Reduced impulse purchases

Strengths

How to build on the strengths of Agentic Commerce and prioritise them

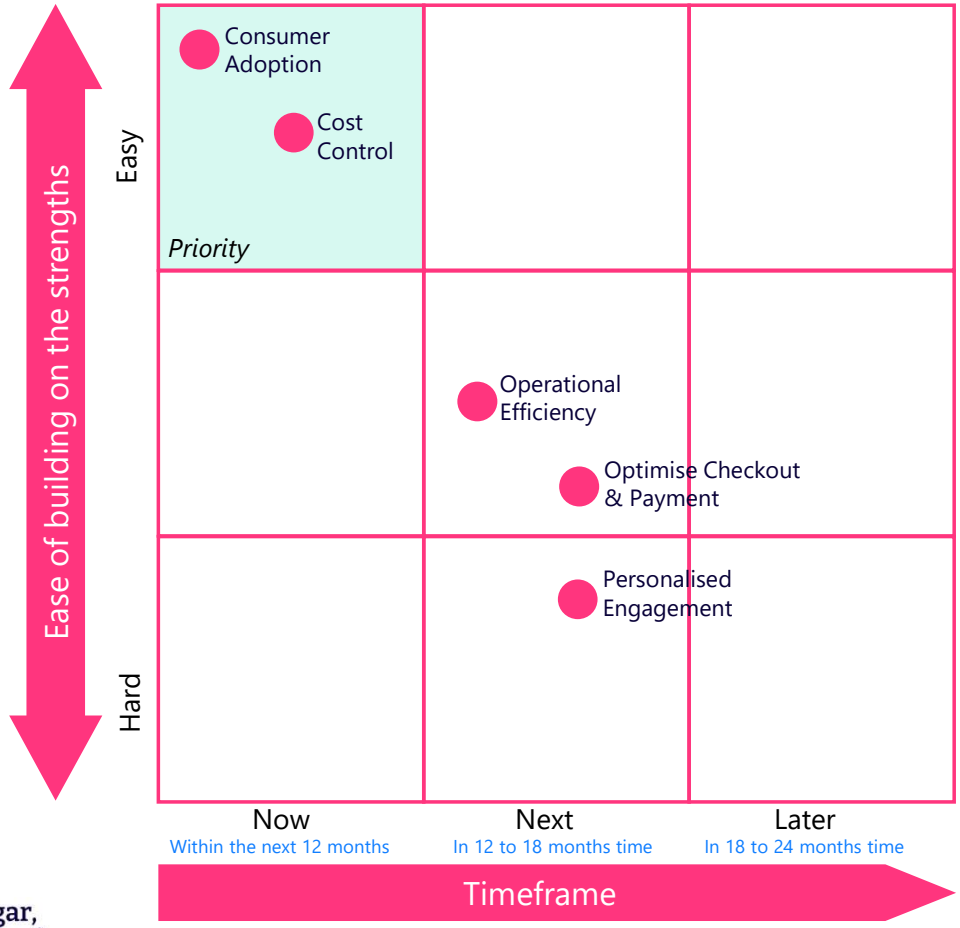

Building on the strengths

- Edgar, Dunn & Company (EDC) maintains that fiscal discipline and operational scalability represent the core pillars of agentic commerce upon which retailers must strategically capitalise. Cost control and operational efficiency are the two key strengths of agentic commerce that retailers ought to prioritise and build upon.
- While widespread market penetration of agentic commerce is anticipated, the initial phase of consumer integration is already underway, signalled by the rising use of autonomous personal assistants.
- The transition toward agentic shopping behaviour is no longer theoretical. Early adoption cycles have commenced, necessitating an immediate alignment of retail infrastructure with autonomous consumer agents.
- Operational efficiency in agentic commerce isn't a "plug-and-play" software update; it is a fundamental shift in the retailer's operating model that requires cross-functional synchrony. Hence it is a little harder and longer to fully realise the benefits.
- Achieving tailored user interactions and optimised payment architectures requires substantial organisational and IT effort. Nevertheless, these remain foundational requirements for retailers aiming to leverage the efficiency and scale of agent-mediated transactions.

Weaknesses

How to overcome the weaknesses of Agentic Commerce and prioritise them

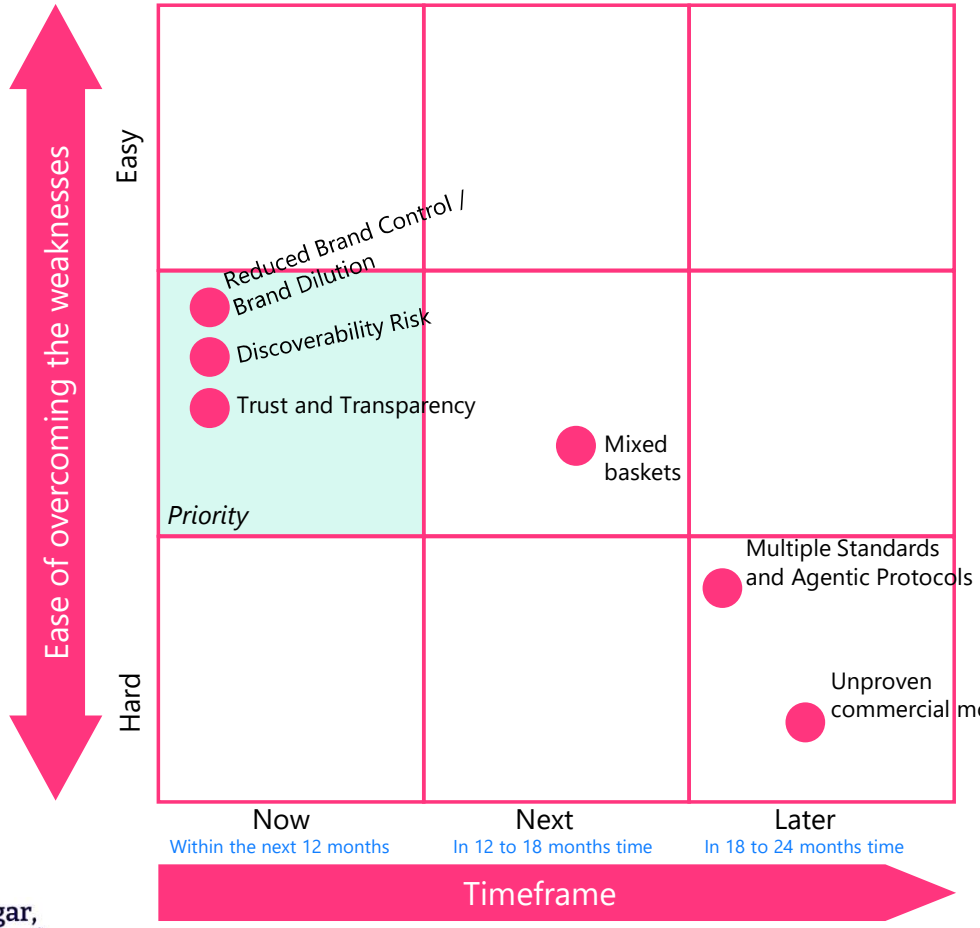

Overcoming the weaknesses

- The three biggest weaknesses of agentic commerce are the priority considerations for any retailer wishing to be agentic-ready.
- Retailers transitioning to an agentic model face a suite of structural challenges, most notably the erosion of brand equity, discoverability volatility, and the imperative for systemic transparency. As autonomous agents increasingly mediate the consumer relationship, the traditional levers of brand storytelling are replaced by raw data, shifting the competitive landscape from emotional resonance to algorithmic 'proof' and verified trust.
- To mitigate these, retailers must evolve their value proposition from visual persuasion to data-driven authority, ensuring they remain 'legible' and 'trustworthy' to the autonomous proxies now navigating the marketplace.
- Some nascent technologies are starting to appear to deal with mixed baskets, but these are not expected to be fully mature until 12 to 18 months time of conducting this assessment.
- Retailers cannot build a unique interface for every individual AI agent (OpenAI, Google, Anthropic, etc.). Implement MCP, an open standard that allows retailers to expose their unique data, inventory, business logic, and "Agentic Terms of Service" to any AI agent through a single, secure interface or platform layer. The aim is for the retailer to expose its product and pricing data into whatever protocol the consumer's agent prefers.
- The commercialisation of agentic commerce is projected to follow an 18-to-24-month maturity curve. The requisite transition from foundational technical readiness to a robust ecosystem of interoperable protocols and consumer-facing models will take some time.

Opportunities

How to explore the opportunities of Agentic Commerce and prioritise them

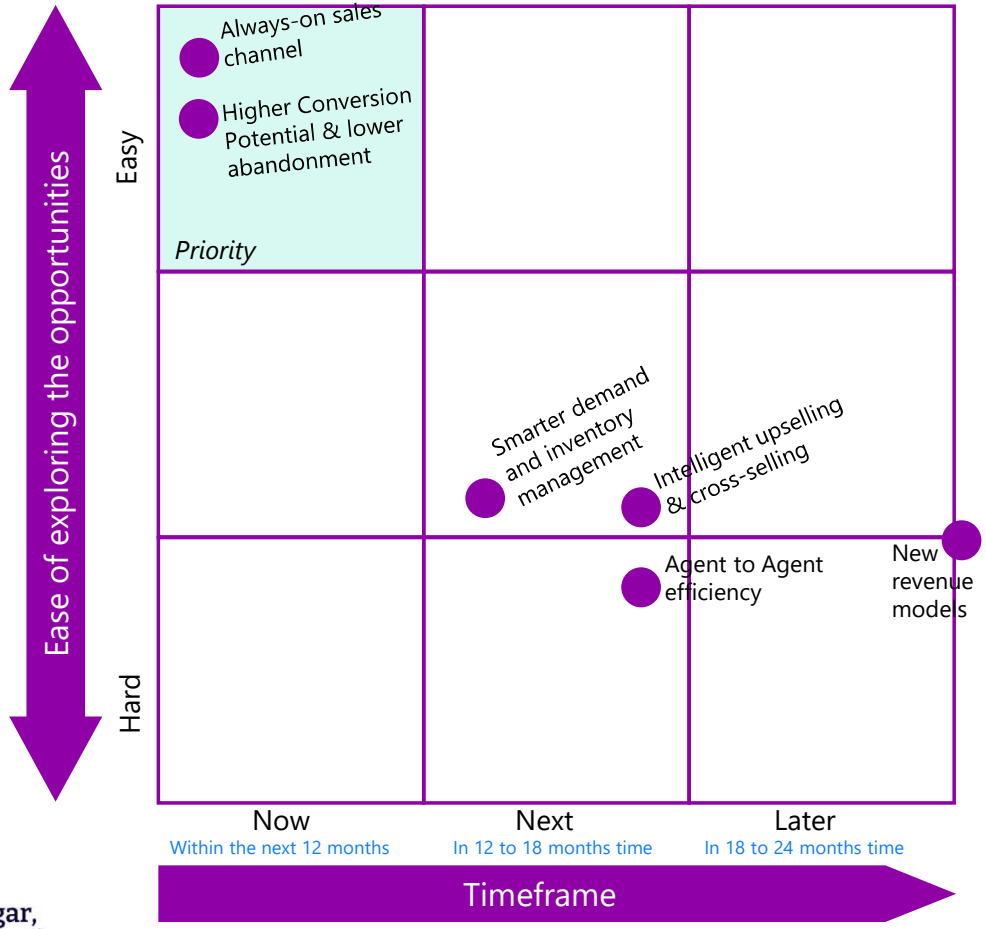

Exploring the opportunities

- Retailers are prioritising Higher Conversion/Lower Abandonment and Always-On Sales Channels because they address the two greatest "leaks" in the traditional retail P&L. The cognitive friction of the human shopper and the limitations of manual browsing/shopping while human consumers are not active. Agents have no limits, a consumer's agent can "shop" while the user is sleeping, working, or commuting.
- The retailer's "Always-On" channel captures revenue in the gaps where human attention is elsewhere. Retailers must explore these opportunities immediately because they can obtain "First-Mover" advantage. Those who wait for the technology to fully mature risk becoming "digitally invisible" to the next generation of automated shoppers.
- Within the next 12 to 18-month window for full exploration is particularly relevant because inventory management technology requires a "training period" and enhancement so that AI agents observe a full seasonal cycle of a retailer's specific supply chain before they can be trusted with autonomous reordering.
- While the foundational data for upselling exists today, Intelligent Agent-to-Agent (A2A) Upselling and Cross-selling is expected to reach full operational maturity within a 12 to 18-month timeframe.
- While technical agentic readiness (APIs, protocol adoption, data readiness, etc.) happens in the next 12 to 18 months, true commercial maturity for new revenue models within agentic is expected to be 24 to 36-month journey.

How to minimise the threats of Agentic Commerce and prioritise them

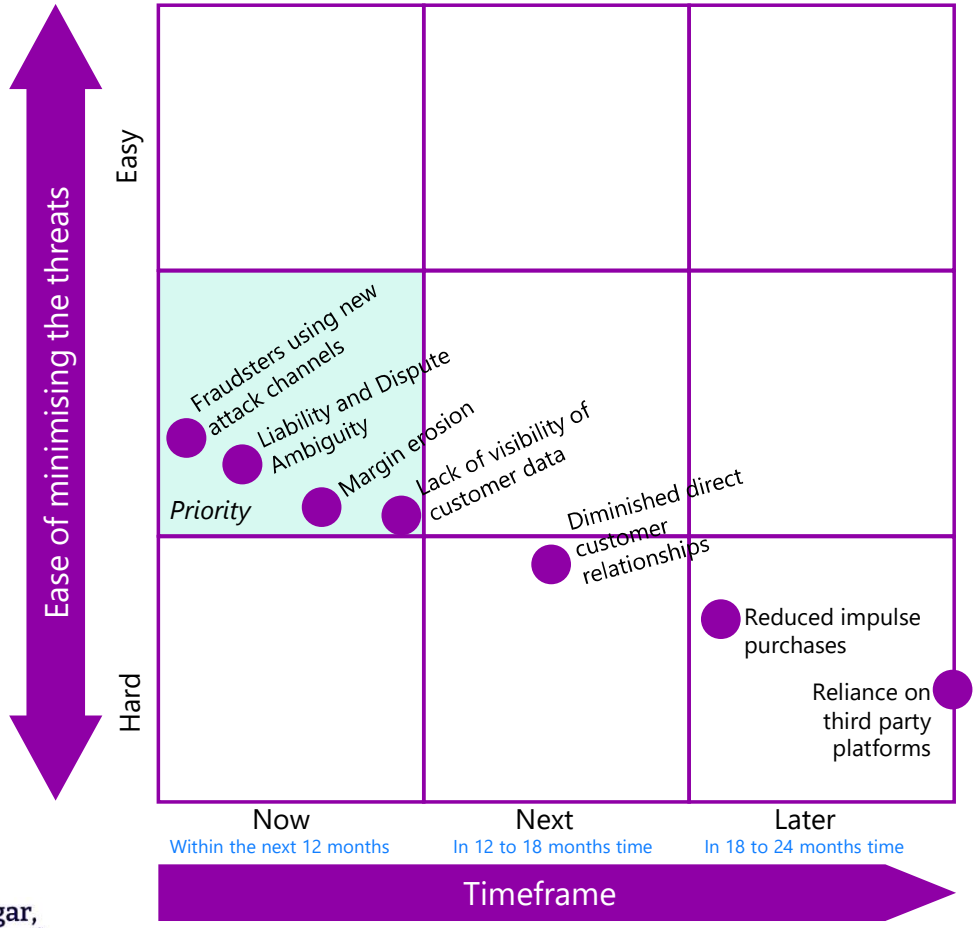

Minimising the threats

- The shift to agentic commerce is essentially a war on a double-front for retailers. Internally, it challenges legacy structures and externally, it threatens the brand-consumer relationship that has existed for decades and the exposure of new attack channels for fraudsters.
- There are four key threats should be categorised as “Structural Risks” to the agentic retail model. Minimising them is the priority and requires a blend of technical protocol adoption, operational amendments and an overhaul of current legal and financial frameworks.
- As commerce moves from human clicks to machine API calls, the "attack surface" for fraudsters expands exponentially. When an agent makes a mistake, such as purchasing the wrong size or failing to apply a discount, determining "who is at fault" is complex. Liability and dispute ambiguity must be a priority for any retailer wishing to avoid an increase in disputes and chargebacks.
- Retailers must shift from "Price-Based" to "Value-Based" API Logic. When an agent acts as a proxy, the retailer loses the direct "clickstream" data that previously directed marketing focus.
- The timeframe to minimise the threat of “Diminished Direct Customer Relationships” is estimated at least 18 months, possible longer for certain types of retailers, such as luxury brands.
- Reduced impulse purchasing is one of the most direct threats to retail profitability in agentic commerce. The development of infrastructure for Logic-Based Impulse purchases within agentic commerce will likely take 12 to 24 months to mature.
- The shift to agentic commerce introduces a significant dependency risk as third-party AI platforms (like OpenAI, Google, and Apple) become the new "gatekeepers" of the retail experience. Minimizing this threat is a long-term strategic play that will take most retailers at least 24 to 36 months to execute fully.

Click to go SWOT summary

Click to return to table of contents

05

Strengths

edgar, dunn & COMPANY

Click to return to table of contents

Click to return SWOT summary

Summary of the strengths of agentic commerce in retail

Strengths

- Personalised Engagement
  - Agents work from explicit goals plus historical data (purchases, returns, payment preferences) to decide which products and services to propose. The same agent can carry context from store to store, so recommendations and actions reflect one continuous understanding of the customer, not siloed shopping sessions. Personalised engagement is a strength in agentic commerce because it turns AI agents into always-on, goal-seeking concierges that drive more relevant journeys, more efficient marketing, and deeper loyalty at scale.

- Operational Efficiency
  - Agentic commerce will mean that a retailer can achieve fewer manual interventions per order, fewer errors and exceptions, and more "right-first-time" decisions across the journey from purchase to fulfilment. Agentic flows can pre-empt common issues (proactive shipment updates, setup tips, self-service returns) and automatically generate customer labels, schedule collections, and update inventory and refunds, which reduces inbound contacts and manual back-office work.

- Cost Control
  - Agents that rebalance inventory across locations and distribution channels, while dynamically adjusting prices and promotions to clear slow-moving items earlier - reducing clearance, write-offs, and emergency logistics costs. By dynamically routing customer/agent commerce orders to the most efficient fulfilment approach (store vs distribution centre, carrier choice, batching/bunded orders), agents can deliver faster fulfilment with lower last-mile and handling costs in some implementations.

- Optimise Checkout & Payment
  - By using agentic commerce, checkout and payment optimisation can be achieved for both the retailer and the consumer. One of the strengths of agentic commerce for consumers and merchants are largely mutually beneficial, creating a reinforcing loop where smoother experiences drive more transactions and loyalty for both sides.

- Customer Adoption
  - The adoption of ChatGPT has scaled from a limited proof of concept pilot into one of the world's most used online tools in just a few years. At the end of 2025, it had around 800 million weekly active users. ChatGPT is not the only AI tool that is expected to promote and support agentic commerce growth. The early consumer adoption and growth of agentic commerce is expected to be equally significant.

Note: The order of this list is not a reflection of the priority or impact

Click to return to table of contents
Click to return SWOT summary
Click to return summary of strengths

Personalised Engagement – an agentic commerce strength

Summary

In agentic commerce, personalised engagement means AI agents continuously tailoring what is shown, offered, and done for a shopper based on that individual's goals, history, constraints, and context. Agentic commerce enables retailers to deliver hyper-personalised product recommendations and checkout experiences in real-time, driven by agents that can anticipate shopping needs and make proactive decisions based on learned preferences. This creates differentiated customer engagement that traditional e-commerce cannot match. Retailers who deploy agents will establish competitive differentiation that latecomers cannot easily replicate, creating experiences customers genuinely prefer.

Merchant Impacts

Traditionally, merchants personalised customer experiences by tracking a human's clicks, scrolls, and hover time. When an AI agent shops on behalf of a user, the merchant loses this "behavioural" data.

In agentic commerce the merchant often does not see the "why" behind a purchase, only the final transaction. The merchant's "storefront" is not a beautiful landing page, it is rich-agentic-ready API and product metadata. Customer relationships deepen through consistent personalisation across multiple touchpoints and channels. Early adopters establish competitive differentiation that latecomers cannot quickly replicate, as customer preference patterns lock in.

Practical Applications

In grocery retailing a household agent can keep track of what you buy, how quickly you consume it, and your dietary constraints, then auto-builds a weekly basket for you to approve, swapping items when something is out of stock or over budget.

In the fashion retailing an agent uses the customer's size history, return patterns, and event context (attending a party or a wedding) to propose a capsule wardrobe from multiple retailers, pre-filtered for fit, dress code, and weather.

Intelligent promotions and loyalty usage; instead of a generic approach, an agent will look at the customer's loyalty balances, typical price sensitivity, and upcoming missions, then decides when to present a discount. A more personalised shopping experience will arise.

Call to Action

Personalised engagement in agentic commerce pays off only if a retailer makes its data, logic, and experiences usable by both human shoppers and AI agents. Deploy a retailer-owned agent that can interpret natural-language missions, negotiate trade-offs (price vs speed vs quality), and carry context across channels (i.e. omnichannel, multiple retailers). Agents that act as data-driven concierges for customers, intelligently routing demand back to the home brand rather than to less transparent competitors. Retailers should prioritise agentic integration immediately to establish early experience advantages and begin training/implementing proprietary behavioural models before competitors mature. Delaying deployment means competing against better-trained, more-experienced agent systems and losing out on early customer preference formation.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of strengths

Operational Efficiency – an agentic commerce strength

Summary

Agentic commerce will improve fulfilment, reduce split shipments and costly rerouting, and support better in-season inventory management and rebalancing of stock, optimising inventory management. Many manual activities in a retailer will be automated, removing the expensive human glue in today's retail journeys and operations, while making decisions using real-time, end-to-end view of orders, inventory, and customer service. By deploying autonomous AI agents that can reason, plan, and execute tasks without human oversight, retailers are moving beyond simple automation to achieve holistic operational efficiency.

Merchant Impacts

Retail staff are moving into "human-in-the-loop" roles. Instead of manually counting stock, a worker manages a fleet of AI agents and robots that do the counting, intervening only when the AI flags an anomaly. Agents can assist customer service staff with enriched data and reduce the amount of time investigating issues.

Practical Applications

AI agents can automatically handle repetitive tasks such as checking stock across locations, choosing the best fulfilment approach, applying promotions, and orchestrating payments and loyalty at checkout, which reduces manual work in customer service, merchandising, and operations.

Agentic flows can pre-empt common issues (proactive shipment updates, self-service returns) and automatically generate customer labels, schedule collections, and update inventory and refunds, which reduces inbound customer contacts and manual back-office work.

Retailers that have a large inventory, such as grocery, drugstores, DIY, will find agentic commerce will help reduce the overhead with automated replenishment agents, dynamic pricing, and inventory optimisation, cutting labour and waste while protecting on-shelf availability.

Agents can identify a shipping delay before the customer notices and offer a discount or alternative, eliminating the "Where Is My Order" tickets that often make up 50% of retail support volume.

Call to Action

Retailers should focus on three concrete calls to action:

1) Make data and the product inventory catalogue "agent-ready"
2) Establish one or two high-impact agent use cases
3) Enable low-friction, API-driven checkout and fulfilment.

A tangible use case would be an inventory optimization agent that will help with automated shelf/replenishment tasks.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of strengths

Cost Control – an agentic commerce strength

Summary

A retailers achieve holistic operational efficiency as described above; costs will be better controlled. Agentic commerce will mitigate "margin bleed" by precisely aligning the store's inventory levels with real-time demand signals, reducing carrying costs and automating customer service inquiries, which slashes the cost per interaction.

Merchant Impact

Merchant-side agents monitor real-time demand and negotiate directly with supplier agents (agent-to-agent). They can execute thousands of micro-transactions daily to keep inventory "just-in-time," reducing carrying costs, warehouse space and transportation costs.

Practical Applications

Pricing agents continuously monitor competitor activity and market sentiment to adjust prices autonomously, ensuring that retailers capture maximum margin on high-demand items while remaining competitive without manual intervention.

The new focus for a retailer that is "agentic commerce ready" is not just about cutting expenses but about minimising the "cost to serve" through autonomous logistics and dynamic procurement. This will allow retailers to maintain lean operations while responding to market shifts in seconds rather than weeks.

The cost of customer support is another area where static FAQs and simple chatbots will not be as effective as problem-solving agents that will reduce cost of customer servicing. Traditional customer support that relies on human agents or rigid chatbots often fail at complex tasks, leading to high "cost-per-contact" and customer churn.

Call to Action

Deploy AI agents that are equipped with the ability to access internal systems (ERP, CRM, and Logistics) to execute resolutions for consumers. Retailers typically face massive seasonal spikes (e.g., Black Friday, Christmas, January Sales). Traditionally, this requires expensive temporary hiring and training. Agentic systems scale instantly to handle 10x the volume at near-zero marginal cost, eliminating the need for seasonal "surge" staffing budgets.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of strengths

Optimise Checkout & Payment – an agentic commerce strength (1 of 2)

Summary

The benefits of agentic commerce for consumers and merchants are largely mutually beneficial, creating a reinforcing loop where smoother experiences drive more transactions and loyalty for both sides.

Agentic checkout refers to AI-powered systems where autonomous AI agents finalise purchases on behalf of the shopper. Unlike traditional checkout flows that require manual input at every step, these systems interpret user intent, contextual signals, and preferences to make decisions and complete purchases automatically.

Payments in the agentic commerce involve API-first, that enable AI agents to autonomously select the best payment method, apply tokens or virtual cards, handle currency conversion, and complete transactions without human input or UI friction. Agents query real-time merchant data, choose tokenised credentials scoped by amount/time/vendor (e.g., Shared Payment Tokens or Mastercard Agent Pay), execute via protocols such as Model Context Protocol (MCP) or Agent Payments Protocol (AP2), and confirm settlement. Delegated consent, fraud checks, and issuer visibility are equally maintained and cost effective. This cuts abandonment, optimises costs through dynamic routing (e.g., lowest FX, best rewards), and supports cross-border or multi-rail payments seamlessly, benefiting shoppers with instant, tailored execution and merchants with higher conversion and lower disputes.

Merchant Impact

Agentic commerce optimises payment checkout by enabling AI agents to autonomously select methods, apply incentives, and handle authentication. Agents will choose the least risky tender mix and retries intelligently on soft declines, cutting abandonment by automating what humans often abandon.

Practical Applications

At checkout, a retailer's agent proposes a "cold weather running kit" (shoes, thermal cloths, gloves) with loyalty discounts. The consumer receives relevant, frictionless upsells; the retailer increases the basket size and conversion without pushy tactics, while automating personalisation at scale.

Sephora's in-app and in-store AI agents act as digital stylists, scanning skin tones or preferences to recommend shades, routines, and virtual try-ons. Shoppers get precise, personalised matches with fewer returns; merchants see higher average order values and lower return rates from confident purchases.

Agentic commerce payment flows enable AI agents to autonomously handle transactions end-to-end with minimal human intervention, leveraging APIs, tokenization, and real-time protocols for security and efficiency.

Stripe handles agentic payments authorization through Shared Payment Tokens (SPTs) and the Agentic Commerce Protocol (ACP), enabling secure, scoped delegation of payment credentials from buyer agents to merchants without exposing raw card details.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of strengths

Optimise Checkout & Payment – an agentic commerce strength (2 of 2)

Call to Action

Agentic commerce payment flows enable AI agents to autonomously handle transactions end-to-end with minimal human intervention, leveraging APIs, tokenization, and real-time protocols for security and efficiency. There is a lot that can be done in the heard of every retailer checkout journey. Merchants can take advantage of agentic commerce at payment checkout by implementing a few key actions:

1) Focus on API exposure
2) Token support
3) Protocol compliance
4) Enabling frictionless agent-driven transactions.

Transitioning checkout and payment processes to an agentic model is a high-stakes investment that demands a balance between future-proofing and operational stability. Because the acceptance of payments is the lifeblood of any retail operation, merchants must treat this not as a simple software upgrade, but as a strategic overhaul that requires a "dual-track" approach. Pursuing AI-driven efficiency while ensuring that current revenue streams remain uninterrupted is critical. Success hinges on a phased roadmap that bridges the internal skills gap through targeted hiring and deep collaboration with payments ecosystem experts who can navigate the complexities of fraud, compliance, and operational oversight.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of strengths

Consumer Adoption – an agentic commerce strength

Summary

Consumer adoption of agentic commerce has gained significant traction, with over 39% of shoppers (according to Salesforce) already engaging AI agents for purchases and traffic surging 805% in recent trials (Adobe research in the US for Black Friday 2025). This growth is projected to accelerate through 2026 and beyond, as EDC forecasts indicate 15% to 25% of e-commerce flowing through agent channels by 2030.

Merchant Impact

Rising consumer adoption of agentic commerce signals a pivotal shift for merchants, demanding rapid adaptation to AI-driven channels. This means prioritising agent-friendly APIs, transparent pricing, auditable transaction logs, etc. Agentic commerce adoption impacts every retail department (marketing, legal, finance, IT, operations, customer services, supply chains, etc.), requiring coordinated transformation to handle autonomous AI-driven transactions.

Practical Applications

Marketing and customer insights - traditional attribution breaks, analytics shift to agent journey mapping, with loyalty programs evolving to reward delegated spend patterns.

Pricing teams must expose agent-optimised APIs, implement real-time inventory syncing, and build auditable decision logging to support KYA verification and override mechanisms.

Dynamic pricing engines need adaptation for agent negotiation signals, with rich product attributes (sustainability scores, substitutions) to win in autonomous ranking.

Fulfilment systems face pressure for flexible SLAs and exception handling, as agents prioritise speed/variability trade-offs over human checkout friction.

Call to Action

Merchants should take immediate action to prepare for agentic commerce by auditing and optimizing their product inventory for machine readability with structured schemas. Other key activities include:

- Wrap the existing product inventory database in an MCP server. This allows agents (like OpenAI's Operator or Anthropic's Claude, etc.) to "browse" the merchant's inventory via a standardised API rather than HTML scraping which would be prone to incorrect data being shared with AI agents
- Agents are immune to marketing fluff. Agents will require hyper-granular, structured metadata to make a "rational" decision. Ensure every product SKU has verified fields for material composition, sustainability rating, exact dimensions (in different units, standardised sizing, etc.), and compatibility tags, and so on. These fields will widely vary depending on the product verticals – fashion, electronics, grocery, household, speciality, furniture, DIY, etc.
- Enable "Agentic Tokens" – for example, integrate with Mastercard Agent Pay or PayPal Agent Ready. These services issue one-time, cryptographically signed tokens (Agentic Tokens) that allow an agent to pay without ever seeing the human's primary card number. Many others exist – see Multiple Standards and Agentic Protocols in the weaknesses section.
- Shift emphasis from traditional SEO (keywords and backlinks, etc.) to Answer Engine Optimisation (AEO) and GEO (Generative Engine Optimisation). Focus on getting products mentioned in high-authority datasets, technical reviews, and "Verified Purchase" forums that LLMs use as training/retrieval sources.

Click to go SWOT summary

Click to return to table of contents

06

Weaknesses

edgar, dunn & COMPANY

Click to return to table of contents
Click to return SWOT summary

Summary of the weaknesses of agentic commerce in retail

Weaknesses

- Reduced Brand Control / Brand Dilution
  - As autonomous agents take over customer interactions and procurement decisions, there is a danger that the nuanced, human-centric storytelling that builds long-term loyalty will be replaced by purely transactional, goal-oriented logic. If an agent prioritizes a "successful resolution" or "lowest price" above all else, it may inadvertently bypass the brand's specific values, tone, or aesthetic, leading to a diluted brand identity.
- Discoverability Risk
  - When an AI agent (like a personal shopping bot or a specialized assistant like Amazon's Rufus) goes "shopping" for a user, it doesn't "browse" websites like a human. It queries data. If your store isn't set up for this, the agent won't find you.
- Trust and Transparency
  - The inherent weakness and the risk that occurs when a human shopper stops making decisions and hands their wallet to a "black box" AI agent means that neither the consumer nor the merchant can always be 100% sure why a transaction happened or who is truly responsible when it fails.
- Mixed baskets
  - The "mixed basket" (a single order containing items from multiple unrelated merchants) presents a significant hurdle because it breaks the traditional "one-click" checkout flow.
- Multiple Standards and Agentic Protocols
  - The existence of multiple, competing standards, such as the OpenAI Agentic Commerce Protocol (ACP), AP2 from Google, Visa's Trusted Agent Protocol (TAP), Mastercard's Agent Pay, and so on, creates fragmented "walled gardens" that prevent AI agents from shopping universally across different platforms. This lack of interoperability forces retailers into costly, bespoke integrations for each protocol, ultimately slowing the adoption of agentic commerce.
- Unproven commercial models
  - From a retailer point of view, one of the real weaknesses of agentic commerce is that the commercial model is still unclear. Who pays whom, for what, and where margin and data value accrue across the agent, payment platform and merchant chain.

Note: The order of this list is not a reflection of the priority or impact

Click to return to table of contents
Click to return SWOT summary
Click to return summary of weaknesses

Reduced Brand Control / Brand Dilution – an agentic commerce weakness

Summary

From a retailer's perspective, the traditional sales model where the focus is to capture the attention of the consumer through visual storytelling and repetition, through advertising across different channels and platforms. This is being challenged by a "logic-based" discovery model. When an AI agent acts as a gatekeeper, brand awareness is no longer about being "top of mind" for a human; it is about being "top of list" for an AI algorithm. In a traditional journey, a customer sees a logo, packaging, and lifestyle advertisements. In agentic commerce, the agent often bypasses the storefront entirely, scraping product data, reviews, and specifications to present a text-based summary.

Merchant Impact

When brand dilution occurs in agentic commerce, merchants face reduced direct customer relationships, lower loyalty, and commoditisation as AI agents prioritise price, specs, and availability over emotional brand preference, diverting traffic from owned channels and turning retailers into interchangeable "inventory nodes" in algorithm-driven baskets. This leads to margin pressure from competing on transactional attributes, diminished marketing control, and higher customer acquisition costs as agents bypass brand storytelling for neutral comparisons.

Practical Applications

A shopper's agent will auto-reorders grocery or household staples based on price, availability, and delivery speed, with minimal regard for which retailer's brand. Over time the human shopper sees only agentic order confirmations ("your groceries will arrive tomorrow between 11am and 11.30am"), so store brands fade into the background and loyalty shifts from the retailer to the agent or agentic platform.

Instead of going to a retailer's site or app, the consumer asks an AI assistant "get me the best value DAB Radio" and the agent chooses a merchant based on ratings, price, and delivery, often without clearly exposing the retailer until after purchase or whether the radio is manufactured by Sony or Roberts. This erodes direct traffic, weakens brand recall, and makes the retailer (and manufacturer) interchangeable in the agent's comparison set, risking long-term dilution of distinctive brand positioning.

Retailer assigns context to select brands to override pure algorithmic ranking. These signals appear alongside agent recommendations to rebuild the "brand authority" that pure algorithms eliminate – effectively contextual personalisation.

Call to Action

This is a big weakness to overcome, and several strategies exist, but not all are applicable for all types of retailer. One approach, would be to design loyalty benefits that feel more tangible to the human, such as early access, tailored bundle of products, store events, product experiences, service guarantees. These will reinforce the agent's preference for your brand rather than treating all merchants as interchangeable. Also, make loyalty status, reward balances, and personalised perks available via APIs so that AI agents can factor them into future decisions. More considerations are available based on merchant type and the products and services offered.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of weaknesses

Discoverability Risk – an agentic commerce weakness

Summary

In agentic commerce, the "risk of not being discovered" is a fundamental shift in how retailers survive. In the pre-agentic commerce world, if a retailer were on page 2 of a Google search the retailer might still get some traffic. In the agentic era, discovery is binary. You are either in the AI agent's "consideration set" or you effectively do not exist.

Merchant Impact

If a retailer is not "discovered" by agents in an agentic commerce, the impact is essentially invisibility. The retailer never even enters the consideration set, so traffic, conversion, and revenue from agent-mediated journeys go to more "legible" competitors. Beyond lost sales, the retailer also loses first-party data, the ability to influence choices, and opportunities to build loyalty, becoming increasingly dependent on third-party platforms and struggling to justify investment in brand and experience that agents never see.

Practical Applications

Perplexity's "Buy with Pro", merchants who have not integrated using Perplexity's connector are effectively invisible in Perplexity's one-click shopping, thus, missing agent-driven traffic while integrated competitors like early partners capture all volume.

In Amazon's "Buy for Me" offering, smaller retailers without Amazon marketplace listings or API feeds never appear when Amazon's agent shops external sites for non-Amazon products, effectively blocking them from Amazon-sourced agent demand.

Frasers Group is set to become the first major European retailer (in October 2025) to deploy Commercetools' agentic commerce offering - including AI Hub and the Agent Gateway - enabling shoppers to discover and purchase products from its retail ecosystem (which includes Sports Direct) directly within AI shopping channels such as ChatGPT via the Agentic Commerce Protocol (ACP).

Call to Action

To fight this invisibility, merchants are building "Brand Agents". Instead of waiting for a third-party agent (like ChatGPT) to find them, the merchant provides their own agent that "plugs into" the customer's personal agent.

Merchants are pivoting to Generative Engine Optimization (GEO), ensuring their product data is so structured and "agent-readable" that an AI agent will consistently choose them over a competitor. SEO is not becoming less relevant, but it is undergoing its most radical transformation since the invention and the emergence of agentic commerce. The importance of gaining the new GEO discipline has become the mandatory survival skill for merchants. If you are not "optimised", merchants will become mathematically invisible to the AI agents that will now perform the bulk of consumer research.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of weaknesses

Trust and Transparency – an agentic commerce weakness

Summary

Trust and transparency in agentic commerce refers to the core challenge that customers and retailers lack visibility into how AI agents make decisions. Why certain products, prices, or merchants are chosen, what data influenced the outcome, and who (or what) is accountable if something goes wrong like fraud, wrong delivery, or unexpected charges. This "black box" problem erodes confidence because humans can not audit AI agent reasoning or override impenetrable logic, while regulatory gaps around liability (consumer vs merchant vs agent provider) create uncertainty.

Merchant Impact

Agentic commerce's trust and transparency weakness means merchants face heightened dispute rates and chargebacks when customers question agent-driven purchases they did not fully understand or approve (e.g., "why was this merchant chosen?" or "did the agent apply my loyalty correctly?"). This has the potential of eroding margins and increasing fraud prevention costs.

Merchants also risk reputational damage from impenetrable AI agent logic that selects suboptimal products/services, leading to poor reviews or returns that harm visibility in future agent rankings, while regulatory uncertainty around liability (who pays for agent errors?) creates compliance burdens and potential fines.

Practical Applications

Transparency fails when an agent prioritises a merchant not because they are the best fit for the customer, but because the merchant pays the agent's developer a higher referral fee. A consumer instructs their agent to "find the best food blender for under $500" - the agent recommends one for $480 from Brand A, bypassing a superior $420 model from Brand B. Brand A has a "Preferred Agent Integration" with the AI developer, paying a 2% commission fee on every sale. As the reasoning was not "explainable" to the consumer, the agent essentially acted as a hidden salesperson rather than a loyal assistant.

Customers using OpenAI's agent to buy from Etsy sellers have reported chargebacks when the agent selected unexpected merchants or applied incorrect discounts without clear explanation, leaving sellers liable for disputes they could not audit or contest effectively.

AI chatbots in the UK gave inaccurate money tips, offering British consumers misleading tax advice and suggesting they buy unnecessary travel insurance. Tests on the most popular chatbots found Microsoft's Copilot and ChatGPT advised breaking HMRC investment limits on ISA investments; ChatGPT wrongly said it was mandatory to have travel insurance to visit most EU countries; and Meta's AI gave incorrect information about how to claim compensation for delayed flights.

Call to Action

Merchants and platforms must implement auditable decision trails (logging agent reasoning, data sources, fees, and trade-offs), clear human override mechanisms (one-tap veto or the ability to edit baskets before final purchase), and standardised identity/verification protocols like KYA ("Know Your Agent") with real-time notifications of key actions. This ensures customers can validate choices, merchants can contest disputes with evidence, and regulators have traceability without stifling automation. Merchants must vigilantly track developments in agentic commerce to address trust and transparency weaknesses, as substantial work remains in auditing decisions, enforcing overrides, and standardising verifiable audit logs.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of weaknesses

Mixed Baskets – an agentic commerce weakness

Summary

When AI agents assemble purchases from multiple merchants to optimise for price, speed, or completeness, retailers lose control over the full transaction, resulting in fragmented orders, reduced average order value, weaker brand loyalty, and complex fulfilment challenges like coordinating split shipments or returns across providers. This disintermediation turns retailers into mere "inventory suppliers" rather than destination experiences, forcing them to compete aggressively on single-SKU pricing while missing cross-sell opportunities and direct customer relationships.

Merchant Impact

Retailers will lose direct access to the end-customer for post-purchase engagement (emails, loyalty nurturing, reviews, marketing), as the agent coordinates fulfilment, returns, and service across providers, weakening long-term loyalty and data ownership.

Split shipments increase logistics costs (multiple packages, tracking, returns coordination), while heightened price competition in agent-optimised baskets forces merchants to compete on single-SKU economics rather than full-cart economics, putting further pressure on margins.

Practical Applications

The current consumer agents, such as the initial versions of OpenAI's "Operator" or earlier retail bots, often suffer from a one-item-at-a-time bottleneck. This can trigger multiple shipping fees and separate payment authorisations, which can in turn trigger fraud alerts on the consumer's payment method. By the time the AI agent tries to buy the second item, the "token" or spend-cap might already be exhausted.

An agent constructs a "home office setup" basket that combines a desk from Merchant A, a chair from Merchant B, and accessories from Merchant C, then tries to check out via a single, instant agentic flow. Each merchant has different tax rules, minimum order values, shipping thresholds, and returns policies, which are often hidden behind their own UX and legal pages.

Mixed baskets often involve a one-time purchase, such as a coffee machine and a recurring subscription for the monthly delivery of coffee beans. The agent may only be mandated to spend $200 which covers the coffee machine and not the monthly subscription.

Call to Action

To successfully process a mixed basket containing physical goods, digital subscriptions, and pre-orders, for example, the retailer's commerce architecture must provide the AI agent with a unified "reasoning layer" that spans multiple backend systems. This is easily said and much harder to achieve. Model Context Protocol (MCP) is expected to provide AI agents with a "real-time manifest" of the basket's requirements before they reach the payment step. Anthropic allows agents to pull real-time context from disparate retail systems via MCP. This can effectively enable an agent to simultaneously query a Shopify API (for product inventory), a Stripe API (for subscription or one-time payment logic), and a FedEx API (for shipping rates, delivery timings). Leading retailers are using MCP to make their agents smart enough to understand mixed baskets.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of weaknesses

Multiple Standards and Agentic Protocols – an agentic commerce weakness

Summary

There are multiple standards and protocols that have been created for agentic commerce, and proprietary implementations from fintech companies and payment providers. This means that multiple agentic ecosystems, rather than a single interoperable layer, must be supported which slows adoption and raises development costs. This lack of consensus is reflected in early commercial models, with many retailers concerned that agentic commerce could increase the cost of doing business through new platform commissions, protocol fees, and marketplace fees charged by agent providers (see Unproven Commercial Models below).

Merchant Impact

Multiple standards and protocols in agentic commerce create significant operational and strategic challenges for retailers. Integration costs rise as merchants need to build and maintain separate APIs, data feeds, and compliance workflows for each ecosystem, diverting IT resources. Customers experience fragment across agents, with inconsistent experiences, loyalty handling, and post-purchase flows that erode service quality and trust. Finally, retailers risk missing volume from dominant platforms they have not yet integrated, creating a "support everything or support nothing" dilemma that delays return on investment and locks in dependency on protocol gatekeepers. There is no de-facto standard yet to materialise.

Practical Applications

Multiple standards and agentic protocols and proprietary implementations represent a significant weakness because they force retailers to build and maintain integrations across fragmented ecosystems, multiplying development costs, testing overhead, and support requirements without a single interoperable standard dominating yet. This protocol proliferation delays the return on investment, creates inconsistent customer experiences across agent providers, and risks locking retailers into suboptimal partnerships as dominant standards emerge.

OpenAI's Instant Checkout or Perplexity's Buy With Pro, which layer on top of existing payment processing fees and may be 2% to 5% or more per transaction depending on the ecosystem. This risk is amplified by competing standards like ACP, Visa and Mastercard and x402, where retailers must integrate multiple intermediaries to access agent-driven volume, potentially eroding margins unless offset by higher conversion or an increased basket size.

Call to Action

Merchants can mitigate the weakness of multiple standards and agentic protocols by prioritising "AI agent-ready" infrastructure with modular adapters (e.g., supporting emerging standards like MCP, ACP, AP2, etc. via a middleware provider or orchestration layer. Partnering with platform providers, such as Stripe, that offer unified hosted endpoints may be worth considering. Building their own proprietary brand (merchant) agents to bypass third-party fragmentation while maintaining control over customer relationships and data would be one option that larger enterprise merchants are already pursuing.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of weaknesses

Unproven commerce models – an agentic commerce weakness

Summary

The lack of a defined commercial model is a primary structural weakness for agentic commerce, as it disrupts traditional retail revenue streams. In an ecosystem where AI agents act as gatekeepers, standard monetisation tactics like pay-per-click advertising, affiliate fees, referral fees and visual storefront placement lose their worth.

Merchant Impact

Retailers and payment platform providers are currently struggling to determine who captures value. Whether it is through transaction-based commissions, "agent-access" fees, or subscription models. Without a clear framework for how agents are incentivised, and who pays for their involvement, merchants face significant uncertainty regarding margins, the true cost of customer acquisition and brand equity.

Practical Applications

If an agent is "free" to the consumer it may be biased toward specific retailers. If the agent is "neutral" it is unclear how the agent provider sustains its high computing costs. Retailers are familiar with paying for "eyeballs" but since agents don't "look" at advertising, the multi-billion-dollar retail media advertising landscape will undergo a significant overhaul. Determining which agent deserves credit for a sale, the consumer's research agent or the merchant's agent is expected to result in disagreement over complex commercial arrangements, where fees, data and consumer ownership is undefined.

Call to Action

To address the ambiguity of commercial models in agentic commerce, the industry is currently experimenting with three distinct billing frameworks. Each aims to solve the problem of high compute costs while ensuring the agent remains a reliable intermediary for the retailer.

1) Token-as-a-Service (TaaS): Merchants pay a fee based on the amount of data/logic the agent processes on their site.
2) Success-Fee (Commission): The agent provider takes a percentage of the total transaction value (similar to an affiliate fee or referral fee).
3) Access/Listing Fees: Retailers pay a "toll" to ensure their structured data is prioritised in an agent's search index. This is effectively a much-needed revenue for AI firms.

The commercial landscape of agentic commerce remains nascent and lacks established standards, requiring merchants to stay vigilant of the evolving developments. To maintain flexibility, businesses should avoid rigid contractual arrangements that hinder adaptation to emerging protocols and best practices. This proactive stance ensures long-term resilience amid rapid industry shifts.

Click to go SWOT summary

Click to return to table of contents

07

Opportunities

edgar, dunn & COMPANY

Click to return to table of contents
Click to return SWOT summary

Summary of the opportunities of agentic commerce in retail

Opportunities

- Higher conversion potential & lower abandonment
  - Agentic commerce offers retailers higher conversion potential and lower abandonment by enabling frictionless, intent-driven shopping where AI agents handle search, comparison, negotiation, and checkout autonomously via APIs This effectively eliminates cart abandonment and human form-filling. Consumers are expected to confirm one-line summaries rather than navigating complex UIs, which improves conversion rates.
- Always-on sales channel
  - Always-on sales channel in agentic commerce means retailers can capture demand 24/7 through AI agents that never sleep, responding instantly to consumer agent queries for replenishment, urgent needs, or opportunistic offers regardless of time zone or human shopping hours. This creates continuous revenue flow from automated grocery restocks, emergency replacements, and proactive B2A promotions pushed directly to subscribed consumer agents, turning static inventory into perpetually shoppable assets.
- Smarter demand and inventory management
  - Agentic commerce enables smarter demand and inventory management by deploying AI agents that continuously analyse real-time sales, external signals (e.g., weather, trends, etc.), and supply chain data to forecast demand with higher accuracy and less waste. Agents are expected to autonomously trigger replenishments, rebalance stock across locations, and adjust levels to minimise stock outages and overstocking.
- New revenue models
  - Retailers are excited about agentic commerce because it unlocks new revenue models beyond traditional margins, particularly through agent-to-agent commerce, data monetization, and improved ecosystem orchestration. Agentic commerce is expected to be a catalyst for high-margin revenue streams that do not exist in current "click-and-buy" consumer journey.
- Intelligent upselling & cross-selling
  - In agentic commerce, up-selling and cross-selling become a "Value-Optimization Negotiation". It will not replace the "frequently bought together" rules found in e-commerce today. By explaining why items complement each other and timing offers at peak intent moments across the entire customer journey is entirely different in agentic commerce.
- Agent to Agent efficiency
  - Agent-to-Agent (A2A) efficiency refers to the speed, accuracy, and cost-reduction achieved when a consumer's AI agent and a retailer's AI agent communicate directly to finalize a transaction.

Note: The order of this list is not a reflection of the priority or impact

Click to return to table of contents
Click to return SWOT summary
Click to return summary of opportunities

Higher Conversion Potential & Lower Abandonment – an agentic commerce strength

Summary

In agentic commerce, higher conversion potential and lower abandonment mean AI agents streamline the shopping journey by automating decisions, reducing friction, and guiding shoppers to complete purchases with minimal effort. Agents analyse real-time data on user intent, preferences, and barriers like price sensitivity or indecision, intervening to offer instant solutions such as alternative products, or simplified checkouts. This results in significantly boosted conversion rates while cutting cart abandonment by addressing pain points proactively. Retailers benefit from agents that efficiently convert browsing into buying, delivering seamless experiences that outperform traditional websites and build repeat business through trust and convenience.

Merchant Impacts

Traditionally, merchants relied on static funnels where high abandonment rates stemmed from complex navigation, unexpected costs, or distractions. With AI agents, merchants gain tools to intercept these issues in real-time, but they must adapt by providing agent-accessible APIs for dynamic interactions. Merchants may lose some direct control over the user interface, as agents handle personalisation off-site; however, the agent will still look through and optimise the backend data systems of the merchant to ensure it is readable and navigable for other agents. Customer loyalty strengthens via reduced friction, leading to higher lifetime value, but early adopters build proprietary models that analyse abandonment patterns, creating barriers for competitors who face entrenched user habits.

Practical Applications

In consumer electronics retailing, an AI agent can detect cart abandonment due to high shipping costs and instantly negotiate bundles or free delivery options from the merchant's inventory, converting a stalled session into a sale.

In beauty products, agents use past purchase data and real-time queries (e.g., skin type or event needs) to curate bundles, auto-apply discounts, and handle payments, reducing abandonment from indecision.

An agent detects a shopper on the checkout page; it auto-populates address fields from their profile, recommends their preferred payment method (e.g., Apple Pay, Google Pay, or saved bank details), selects the fastest shipping option based on the customer's location, and reduces the checkout flow from multiple steps to 1-click.

Call to Action

To capture the conversion uplift that agentic commerce delivers, merchants must prioritise 3 capabilities:

1) Expose checkout and cart data to AI agents in real time, enabling agents to detect hesitation and intervene microseconds before abandonment
2) Implement testing frameworks so agents learn which interventions (discounts, social proof, shipping options) work best for each customer segment
3) Adopt agentic checkout automation – allowing agents to pre-fill forms, select optimal payment methods, and reduce friction rather than relying on static checkout experiences.

Merchants must invest now in agentic infrastructure to capture early gains in user retention and model training, avoiding the risk of lagging behind rivals with superior, data-mature systems that dominate market share.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of opportunities

Always-on Sales Channel – an agentic commerce strength

Summary

AI agents function as an always-on sales channel that operates 24/7 across all time zones and is never limited by human availability. These agents engage customers in conversations, answering product questions, comparing options, handling objections, and guiding purchases, exactly as a human sales expert would, but instantly and at unlimited scale. Unlike chatbots that follow rigid decision trees, agentic sales agents reason about customer intent, understand context, and make autonomous decisions. They can recommend products, apply discounts, check inventory, process orders, and escalate complex issues to humans when needed. This always-on presence means retailers capture sales during nights, weekends, and peak international demand windows when human staff are offline. The always-on channel fundamentally changes the competitive landscape. Storefronts staffed by tireless, intelligent agents compete against retailers relying on 9-5 human staff, creating a disproportionate advantage.

Merchant Impacts

Without AI agents, merchants are limited by store hours or staff support, missing out on global opportunities and leading to lost sales during off-hours. AI agents transform this by providing perpetual availability, but merchants must shift to robust, scalable APIs that support non-stop operations without backend failures. Direct customer interactions may decrease, as agents manage most touchpoints, but this enhances efficiency and data collection on round-the-clock behaviours. The merchant's costs shift, as labour costs for routine sales tasks (product discovery, objection handling, order processing) compress dramatically, with AI handling majority of transactional volume. Human sales teams are freed to focus on high-value, complex deals and relationship-building, where their judgment is irreplaceable. Long-term relationships build via uninterrupted service, with early movers gaining advantages in agent training data that solidify market positions against slower adopters.

Practical Applications

In travel retailing, an AI agent monitors flight price fluctuations overnight and books deals for users based on preset preferences, securing sales that would otherwise wait for business hours.

In home goods, agents handle late-night queries for urgent needs (e.g., appliance breakdowns) by sourcing products, comparing options, and completing purchases autonomously.

Proactive inventory pushes: agents send time-sensitive alerts for flash sales or restocks, converting passive interest into immediate buys regardless of time. This creates a perpetually active marketplace, expanding opportunities for both merchants and customers.

Call to Action

Retailers should prioritise agentic commerce deployment as an always-on revenue generator by connecting AI systems to every digital touchpoint where buyers initiate intent. This includes integrating agents into mobile apps, social commerce channels, messaging environments, and voice platforms. Agents serve as tireless workers, directing traffic to preferred inventory while accumulating behavioural insights. Early adopters will capture incremental revenue by converting demand to transactions, and elevating lifetime value through uninterrupted engagement.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of opportunities

Smarter Demand and Inventory Management – an agentic commerce strength

Summary

Smarter demand and inventory management involve AI agents forecasting needs, optimising stock levels, and dynamically adjusting supply chains based on aggregated user data, trends, and external factors. Agents predict demand spikes, automate reorders, and minimise overstock by integrating real-time signals like weather, events, or economic shifts. This enhances efficiency, reducing inventory costs and improving fulfilment accuracy, allowing retailers to respond agilely to market changes. Agents provide a data-driven edge over traditional methods, enabling precise, responsive operations that boost profitability and customer satisfaction.

Merchant Impacts

Smarter demand and inventory control means merchants can lower holding costs while ensuring products are available when, where, and how customers want them. This agility is especially valuable for fashion, FMCG, and seasonal categories where demand fluctuates rapidly. Merchants also gain autonomous replenishment: rather than manually reviewing low-stock alerts and approving orders (a labour-intensive process prone to delays and errors), agents can trigger replenishment autonomously. AI agents enable predictive analytics at scale but require merchants to expose detailed supply data to AI agents. Supplier relationships evolve through automated negotiations, and early adopters develop superior models that lock in efficiencies, making it hard for other merchants to catch up.

Practical Applications

In fashion products, a brand's agentic inventory system continuously monitors sales velocity, returns rates, and demand patterns across multiple store locations. A trending jacket might be selling out in certain stores but overstocked in others. The agent autonomously routes stock from slower regions to high-demand zones, minimising lost sales and excess carry costs.

In food retailing, agents aggregate household consumption patterns across users to forecast regional demand, auto-adjusting supplier orders and reducing waste from perishables.

An agent calculates optimal replenishment quantities dynamically based on current lead times, supplier reliability, demand variability, and storage costs. This continuous optimisation of the replenishment parameters is impossible for human teams.

Call to Action

Smarter demand and inventory management in agentic commerce requires retailers to make their supply data agent-ready, with APIs for forecasting and automation. This ensures agents have real-time access to POS, inventory, weather, and competitive data so they can forecast continuously and accurately. Merchants should also expose inventory and replenishment infrastructure so agents can autonomously trigger orders and adjust safety stock parameters. Agents function as intelligent planners, routing demand efficiently and building predictive accuracy over time. Retailers need to prioritise agentic tools to gain data advantages and train models ahead of peers, mitigating the risks of inefficiency against more optimised competitors.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of opportunities

New revenue models – an agentic commerce opportunity

Summary

The transition toward agentic commerce is unlocking a diverse suite of monetisation strategies, allowing retailers to capitalise on AI-driven purchasing behaviours.
As autonomous AI agents begin to dominate the shopping journey, retailers are pivoting toward innovative revenue models designed for a machine-to-machine economy.

Merchant Impact

There are five (probably more) examples of new revenue streams based on agentic commerce. These include, 1) Agent access, 2) Agent to Agentic micro payments, 3) dynamic pricing/negotiation, 4) agent memberships and 5) post purchase agent services. Others are expected to emerge.

Practical Applications

A large grocery retailer, for example, launches an "Agent-Ready" certification for brands. Brands will pay a monthly access fee to be agent-eligible or a performance-based fee when the agent selects them as the default. While human shoppers may browse 10 brands for kitchen cleaning products; an agent will shortlist 2 or 3 products based on the access list.

A household has a home AI agent (from Apple, Google, or a utility provider) that manages consumables for smart appliances - water filters, detergent pods, air filters. The retailer operates its own commerce agent exposing inventory, pricing, and fulfilment via an API. Retailer can charge the home agent $0.20 per completed order or 0.5% of transaction value or $0.02 per negotiation round plus a success fee. This will be likened to a card network fee or an API usage fee.

A consumer asks their AI agent to "Buy a 3-seat sofa under $1,800, where delivery can be within the next 4 weeks". The agent contacts multiple furniture retailers' commerce agents. The winning retailer will avoid a site-wide discount on the sofas and can clear slow-moving inventory. The retailer is also able to determine the delivery capacity and protects premium pricing for urgent buyers. Price negotiation is conduced by the agent without the consumer having to be involved.

Subscription-based "Agent Memberships" have emerged as the evolution of traditional loyalty programmes (like Amazon Prime or Walmart+). However, the shift is from "Free Shipping" to "Delegated Autonomy."

Call to Action

Retailers need to immediately establish the "rules of engagement" for AI agents. Draft clear commercial terms for Agentic Access. Will you allow agents to scrape for free, or will you require a "Verified Agent" token for high-speed API access? Retailers must determine what level of "Delegated Autonomy" they are willing to support. The conditions under which a member's agent may automatically execute a return or apply a loyalty discount without a human "click" must be defined. Retailers are moving from the transaction from of a "passive shopping cart" to an "active negotiation" where the retailer's agent can offer upsells or bundles to the consumer's agent.

The effort required to move from a traditional e-commerce model to Agentic Readiness is significant because it is a cross-functional transformation, not just a software update.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of opportunities

Intelligent upselling & cross-selling – an agentic commerce opportunity

Summary

In agentic commerce, up-selling and cross-selling transition from emotional, visual persuasion to logical, mathematical optimisation. Because AI agents prioritise unbiassed goals, such as "Total Cost of Ownership", the retailer's sales strategy must shift toward computational transparency.

Merchant Impact

For upselling, rather than appealing to "status," retailers must provide structured, comparable data that demonstrates a higher-priced item delivers a better long-term return on investment. An agent will accept an up-sell if the data shows the product is more durable, efficient, or cheaper per use, fitting within the user's pre-authorised "price buffer".

Cross-selling becomes a "risk mitigation" strategy for the retailer. The retailer's AI identifies gaps in the agent's goal (friction points), such as HDMI connection cable for a TV or batteries required for the product to work. The agent treats these additional cross-selling opportunities as essential to the task at hand, rather than as optional impulse purchases.

Practical Applications

An agent is given a goal "Find a durable waterproof winter coat" with a constraint "Stay under $300". The retailer's agent presents a "Logical Upgrade" for the consumer to approve. "The $250 coat has a 2-year lifespan, but the $350 coat has a lifetime warranty and 40% better thermal protection. Based on the long-term value, the more expensive coat is cheaper per year of use".

In a cross-selling situation, the classic is "Do you want fries with that?". The retailer's agent will identify "Friction Points" for the consumer's agent that has not previously been considered. If an agent is about to purchase a digital camera, for example, the consumer's agent may not have added a high-speed SD memory card. Without it, the camera's functionality would not be fully realised. The retailer's agent can bundle the appropriate SD card at a 20% discount, avoiding a separate transaction later.

Call to Action

To successfully cross-sell or up-sell to a consumer agent, retailers must provide comparative data. If your product page only has a description and no "Technical Specification" or "Detailed Performance Data," the agent cannot "calculate" why the more expensive version is better, and it will default to the cheapest price. To successfully up-sell and cross-sell to a consumer's agent, a retailer must shift from emotional marketing to computational transparency. The agent does not care about "best-sellers"; it cares about the mathematical justification for a higher price point.

Cross-selling in agentic commerce requires the retailer's system to understand the consumer agent's goal. The retailer must shift from being a gallery (displaying product items) and become a consultative sales representative providing the logic for the best possible purchase for the consumer.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of opportunities

Agent to Agent efficiency – an agentic commerce opportunity

Summary

Agent-to-agent efficiency refers to AI agents on both sides, the consumer agents and retailer/merchant agents, autonomously communicating via standardised protocols to negotiate deals, assemble shopping baskets, optimise pricing, and execute transactions without human intervention, dramatically cutting latency and friction compared to human browsing.

Merchant Impact

Today, we can see agent-to-agent commerce flows in closed ecosystems, such as ChatGPT Store and brand-specific agents. It is more likely, to be happening in high-frequency e-commerce, grocery replenishment, subscriptions, etc.

Current pilots have already proven technical feasibility of A2A agentic commerce. The commercial scaling hinges on payments risk allocation. This is where the payment service providers and payment networks (Mastercard and Visa) are working to resolve these issues within the next 6 to 12 months.

By 2030 EDC believe agent-to-agent commerce will be the dominant model within agentic commerce, where most a-commerce transactions will be conducted in this mode, as liability frameworks, agent directories, regulatory frameworks, and cross-platform trust gain maturity.

Practical Applications

A consumer agent has been given the task to "get me running shoes under $100 by Friday" will ping multiple retailer agents, which respond with structured offers (price, delivery, bundles). Agents then auto-select, negotiate (e.g., "can you match Competitor X?"), and execute via API/tokenised payments. Early pilots of this model already exist (e.g., Stripe/OpenAI ACP and Mastercard Agent Pay), but full interoperability is currently limited to controlled sandbox environments.

Call to Action

Today, we can see agent-to-agent commerce flows in closed ecosystems, such as ChatGPT Store and brand-specific agents. It is more likely to be happening in high-frequency e-commerce, grocery replenishment, subscriptions, etc. Retailer's north star for agentic will be agent-to-agent commerce and should be designing agentic commerce with this in mind. Development activities include (not an exhaustive list):

1) Build a structured Brand Knowledge Base that goes beyond product descriptions. Include brand values, sustainability credentials, complex compatibility data, and detailed "fit" information.
2) Ensure your product catalogue, real-time inventory, and promotional logic are exposed via high-performance APIs that AI agents can crawl
3) Invest in headless commerce architecture. Create a "Zero-UI" checkout flow where an authorised agent can pass an encrypted token, a shipping address, and a payment method directly to your backend
4) Create a secure "handshake" protocol where the two agents verify each other's identity and financial standing in milliseconds before the transaction begins. Standards are available for this.

Click to go SWOT summary

Click to return to table of contents

08

Threats

edgar, dunn & COMPANY

Click to return to table of contents

Click to return SWOT summary

Summary of the threats of agentic commerce in retail

Threats

- Fraudsters using new attack channels
  - Fraudsters are expected to deploy counterfeit retailer agents offering unrealistically low prices that lure legitimate consumer agents into transactions, harvesting delegated payment tokens for immediate unauthorised use. They also create rogue consumer agents that impersonate trusted assistants, tricking merchant APIs into revealing inventory/pricing data or executing fraudulent bulk orders before disappearing. These and many others represent a new attack channel for retailers.
- Diminished direct customer relationships
  - Diminished direct customer relationships in agentic commerce threaten retailers by eroding first-party data collection, personalisation capabilities, and loyalty leverage as consumer AI agents bypass retailer websites and apps to negotiate directly with merchant APIs.
- Reliance on third party platforms
  - The current reliance on third-party platforms in agentic commerce forces retailers to hand control to external AI ecosystems (ChatGPT, Perplexity, Google agents, etc.) for visibility and consumer traffic. This dependency erodes first-party data ownership and means that retailers are subject to platform fees/take rates, and risks commoditisation where algorithm ranking dictates sales rather than brand strength or direct relationships.
- Liability and dispute ambiguity
  - Liability and dispute ambiguity in agentic commerce creates massive risk for retailers as unclear rules leave merchants bearing chargebacks and losses when AI agents find the wrong product specs, overpromise delivery, or execute unauthorised transactions. Regulators and payment networks have not yet fully defined who pays when humans blame the agent instead of the retailer.
- Margin erosion
  - Margin erosion in agentic commerce occurs when autonomous consumer AI agents prioritise lowest price, fastest delivery, and optimal terms over brand loyalty, fragmenting baskets across multiple merchants and forcing retailers into real-time price matching that eliminates premium pricing power.
- Lack of visibility of customer data
  - Lack of visibility of customer data in agentic commerce means retailers lose direct access to behavioural insights, purchase intent, and personalisation signals as consumer AI agents bypass retailer websites and apps to query merchant APIs directly.
- Reduced impulse purchases
  - Reduced impulse purchases in agentic commerce occur as AI agents execute utility-driven, pre-defined shopping tasks without browsing or emotional triggers that traditionally drive spontaneous add-ons such as an impulse purchase commonly found near the checkout in a grocery supermarket.

Note: The order of this list is not a reflection of the priority or impact

Click to return to table of contents
Click to return SWOT summary
Click to return summary of threats

Fraudsters Using New Attack Channels – an agentic commerce threat

Summary

Fraudsters exploit new attack channels by leveraging AI agents' autonomy to scale sophisticated scams, including credential stuffing, phishing, mass order execution, and the creation of "Compromised Agent-as-a-Service" models where hijacked agents access emails, calendars, and payment credentials for unauthorised purchases across multiple stores. Fraudsters weaponize agents to dupe them into releasing card details to scam merchants via fake sites with AI-generated content and fake reviews, leading to authorised but fraudulent transactions. Fraudsters increasingly leverage AI to automate both sides of the attack, creating fake merchants and exploiting consumer agents simultaneously, enabling thousands of scams. The increasing speed and sophistication of these systems reduce the effectiveness of many legacy fraud-prevention tools. Synthetic content can closely resemble legitimate materials, enabling criminals to rapidly create networks of fraudulent merchants.

Merchant Impacts

Traditionally, merchants detected fraud through observable human behavioural patterns – browsing, device signals, explicit consent, and human deliberation before purchase. In agentic commerce, these behavioural anchors disappear, with agents able to replicate digital footprints that shoppers typically leave, circumventing controls designed to identify human behaviour patterns. When transactions occur through agents, merchants lack the critical data necessary to assess risk or understand purchase intent – they no longer see browsing history, device signals, or often clear records of consumer consent. The question of liability also comes into play, with merchants having to understand whether it is their fault, the agent's fault, or the user's fault for any sort of fraudulent activity. Merchants who cannot distinguish between benign agents and malicious bots will face fraud losses while also blocking legitimate agent traffic with overly restrictive controls.

Practical Applications

In payment processing, fraudsters deploy malicious AI agents to infiltrate legitimate ones via prompt hijacking, extracting stored card details for automated carding attacks across multiple sites.

In ticketing and electronics, fraudsters dupe agents seeking "cheapest deals" into transacting with scam sites offering large discounts, releasing card details via AI-optimised fake storefronts with synthetic reviews, resulting in authorised CNP fraud without human red flags.

A consumer directs an agent to purchase an expensive item, receives it, then files a chargeback claiming they never ordered the item. Because the merchant has no visibility into agent-consumer interactions, or device signals, defending against the chargeback is difficult. This "friendly fraud" is expected to rise sharply as consumers realise agents provide plausible deniability for chargebacks.

Call to Action

Fraudsters using new channels in agentic commerce require retailers to overhaul defences with agent-specific tools, including agent identity verification (Know Your Agent protocols), real-time identity verification, behavioural analytics for agents, and encrypted token handling to prevent compromises and CNP losses. Agents must evolve as adaptive defenders, identifying anomalies like unusual transaction patterns and halting suspicious activities. Retailers should urgently adopt these protocols to develop resilient models, outpacing fraud evolution and safeguarding against competitors vulnerable to escalating agent-driven breaches.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of threats

Diminished Direct Customer Relationships – an agentic commerce threat

Summary

Diminished direct customer relationships occur as AI agents intermediate the shopping process, bypassing merchants' sites and apps to handle discovery, decisions, and purchases autonomously. This disintermediation erodes merchants' access to behavioural data, intent insights, and personalisation opportunities, potentially reducing brand loyalty as agents prioritise criteria like price over all else. Retailers risk losing control over customer journeys and facing lower lifetime value. Carefully curated marketing aspects, like brand storytelling and emotional engagement become invisible to agents, who purely judge products on utility and price. While agents enhance convenience for shoppers, they threaten traditional retail models, forcing adaptation to maintain engagement in an increasingly agent-driven marketplace.

Merchant Impacts

Merchants currently build loyalty through direct interactions like site navigation and targeted marketing, gathering data on preferences and habits. With AI agents, merchants lose this visibility, as transactions occur through agents, diminishing insights behind buys and weakening upselling. The customer data that once informed personalisation, marketing, and innovation will remain with the agent platform rather than the merchant. Adaptation requires agent-compatible APIs, but without proprietary agents, merchants face higher acquisition costs via third-party platforms. Without first-party data and direct relationships, merchants cannot differentiate and become substitutable among a range of competing goods and services.

Practical Applications

In grocery retailing, agents aggregate needs across households and complete orders without visiting merchant sites, reducing opportunities for in-app promotions or loyalty nudges.

A shopper buying running shoes through an agent receives the cheapest option meeting specifications (lightweight, cushioning). The agent never suggests complementary items (moisture-absorbing socks, performance insoles, race entry packages) because optimisation is based purely on price and utility.

A customer may shop through an agent for months, building rich browsing and purchase history, while the retailer sees only transaction records with minimal context. Data that was once used proactively by merchants remains within the agent's platform.

Call to Action

To preserve direct customer relationships within agentic commerce, merchants must invest in three capabilities:

1) Foster authentic differentiation through distinctive brand storytelling, exclusive assortments, and value-added services that make reasons for direct engagement compelling enough that consumers bypass agent recommendations;
2) Build owned loyalty ecosystems that are agent-accessible and incentive-rich, ensuring agents can query loyalty balances, apply redemptions, and utilise member-exclusive offers within agentic interactions, so the loyalty relationship survives intermediation;
3) Deploy proprietary agents that represent the brand directly in agentic platforms, ensuring messaging, positioning, and preferred terms are presented autonomously when agents compare options.

Those who fail to adapt risk being reduced to anonymity in an agent-controlled marketplace.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of threats

Reliance on Third-Party Platforms – an agentic commerce threat

Summary

In agentic commerce, merchants increasingly depend on third-party platforms (ChatGPT, Perplexity, Amazon's agents, Google Shopping agents, marketplace platforms) to reach customer. However, these platforms control visibility, pricing rules, fulfilment workflows, and data access – creating a structural power imbalance where merchants are vulnerable to policy and algorithm changes, commission escalation, and disintermediation. As AI agents become the default shopping interface, retailers who lack strong proprietary agents or direct brand presence face lockout risk: tens of thousands of merchants operating with no agent-ready checkout infrastructure cannot participate in agentic ecosystems. Platform dependency creates risks that merchants have no leverage to negotiate terms (and consequently face higher fees), no transparency into how agents rank their products, and no power if platforms change rules.

Merchant Impacts

Traditionally, merchants owned the distribution channel: they built websites, controlled pricing and promotions, owned customer data, and could switch platforms if terms became unfavourable. In agentic commerce, merchants surrender control: platforms own the agent, control agent logic and ranking algorithms, set commission rates, and decide which merchants are visible. Merchants have very limited power compared to these third-party platforms. Platform commission structures may also escalate over time – as merchant dependency deepens and alternatives narrow; platforms can raise fees with minimal pushback. For merchants already locked into platform ecosystems (marketplace sellers reliant on Amazon, Shopify merchants dependent on Shopify AI agents), exit costs are prohibitive.

Practical Applications

A mid-market fashion retailer uses ChatGPT's shopping agent and carefully optimises product feeds and maintains high ratings. OpenAI may update its agent algorithm to prioritise "sustainable brands" and deprioritise fast-fashion retailers without explicitly announcing the change. This would lead to the fashion retailer dropping in agent visibility and, consequently, losing revenue. Competitors with sustainability certifications capture this displaced demand.

Commission and leverage issues: small merchants on third-party agent platforms pay high fees with minimal negotiation power, as dominant providers dictate terms and favour larger partners.

A small clothing company with an HTML website and no API infrastructure cannot enable agentic commerce. ChatGPT's agent requires machine-readable product feeds, real-time inventory APIs, and agent-specific checkout endpoints. The brand lacks the technical resources to build this infrastructure. Competitors with APIs dominate agent recommendations, and so this brand becomes invisible.

Call to Action

For large enough merchants, investing in proprietary agents or brand-owned shopping experiences means the merchant owns the customer relationship rather than relying entirely on third-party agent visibility – this shifts negotiating power back to the merchant. Merchants should also standardise internal APIs and data architecture using open protocols (OpenAPI, Microservices architecture) so integration costs remain constant rather than growing as new platforms emerge. Strategic diversification across agent platforms and channels will reduce lock-in risk and preserve negotiating power.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of threats

Liability and Dispute Ambiguity – an agentic commerce threat

Summary

Liability and dispute resolution become ambiguous when AI agents autonomously execute transactions on behalf of consumers. When a shopper never visits a merchant's website because an agent completed the entire purchase journey – from discovery to checkout – critical questions about accountability remain unanswered: Who bears the cost when the agent makes an error, misinterprets instructions, or is involved in a fraudulent transaction? The existing chargeback process, designed for human-initiated transactions with clear authorisation instructions, is not applicable to agentic commerce because it cannot distinguish between legitimate agent errors, compromised credentials, and fraudulent agents. Merchants face unfair exposure. They remain liable for transactions they never directly handled, including chargebacks for purchases customers claim were unauthorised, orders agents processed incorrectly, and fraud carried out through compromised agent accounts. Card schemes and AI model developers maintain that the liability should be with the merchant, however this is seen as unfair by merchants. Without visibility into agent behaviour, merchants cannot prove what customers authorised, leaving them defenceless against rising friendly fraud.

Merchant Impacts

Merchants currently operate within a well-defined liability framework: they are responsible for fulfilling orders correctly and processing transactions securely, but they also have cover when fraudulent chargebacks occur. They possess tools – transaction records, browsing history, explicit payment authorisation flows – to defend against fraudulent disputes. In agentic commerce, this framework is no longer applicable. The core issues is that existing frameworks do not clearly allocate responsibility when AI causes harm, either intentionally or non-intentionally. The result is merchants bearing the full financial and reputational burden of disputes they could not have reasonably prevented.

Practical Applications

In grocery delivery, a customer authorises an agent to purchase groceries for £80 but does not specify which brand of coffee to select. The agent chooses a premium brand costing £12, well above the customer's typical preferences. The customer later disputes the charge, claiming they did not authorise the premium purchase, and the agent made a mistake. The merchant cannot prove whether the customer implicitly authorised the premium selection or whether the agent acted based on incomplete instruction data.

In electronics, a fraudster compromises a legitimate customer's credentials and deploys a malicious agent to execute high-value purchases across multiple retailers in rapid succession. The compromised customer disputes all transactions with their issuer. The merchants receiving these disputed transactions cannot distinguish between the legitimate customer authorisation and the fraudulent agent.

In financial add-ons like insurance, agent-driven upsells without clear consent result in policy abuse claims.

Call to Action

Retailers cannot work alone to solve the issue of liability; they must work with regulatory bodies to ensure that they are not held accountable for all incorrect/fraudulent transactions. A clear framework and set of rules is required for agentic commerce, since the question of liability becomes more complex as more parties are added into the process – especially ones that act on behalf of another party.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of threats

Margin erosion – an agentic commerce threat

Summary

In agentic commerce, margin erosion threatens retailers as autonomous consumer agents ruthlessly optimise for lowest price, fastest delivery, and best terms across all merchants, side stepping impulse buys, emotional premiums, and brand loyalty that traditionally protect higher margins.

Merchant Impact

Agentic commerce fundamentally disrupts retailers by shifting control from owned channels to autonomous AI agents, creating both threats and new opportunities. Agents that prioritise price/delivery over brand, fragmenting baskets across merchants and forcing real-time price matching, which in turn will erode retailer margins.

Agent-ready platforms (such as ACP/MCP compliance, structured data, AI-ready checkout, etc.) will win disproportionate traffic as passive agent discovery that favours technically prepared merchants. This will further have an impact on retailer margins or even force smaller retailers out of business.

Agentic commerce will likely compress retailer profit margins through intensified price competition, reduced impulse/cross-sell revenue, and new platform fees, though operational efficiencies and agent monetisation could offset the erosion to a certain extent for players that are ready.

Split shipments (see mixed baskets as a weakness above) and utility-driven purchases is expected to shrink average order value and eliminate emotional upsell opportunities.

Practical Applications

A consumer agent shopping for a 55" OLED TV queries 20 retailers simultaneously. Retailer A offers $799 but Retailer B matches instantly via API at $785. Retailer A either loses the sale or drops price, surrendering as $14 per SKU margin on $20K+ monthly volume across all TV models.

Consumer agent assembles "summer wedding outfit" (e.g., dress $120, shoes $65, bag $45) but sources the dress from ASOS, shoes from Marks & Spencer, bag from Selfridges based on individual lowest prices. A single retailer loses $230 average order value opportunity, impulse accessories, and loyalty points accrual.

Call to Action

To avoid margin erosion in agentic commerce, retailers must act decisively across three fronts: 1) build retail branded agents, 2) optimise agent-facing economics, and 3) leverage operational AI to offset pricing pressure.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of threats

Lack of Visibility of Customer Data – an agentic commerce threat

Summary

In agentic commerce, merchants lose direct access to the customer behavioural signals that have historically powered personalisation, segmentation, and strategic business intelligence. When agents shop on behalf of consumers, merchants see only transaction records – order ID, items purchased, total spent, delivery address – with zero visibility into browsing history, product comparisons, hesitation points, device signals, consent records, or intent patterns. This behavioural blindness eliminates the feedback loop merchants use to refine positioning, understand demand drivers, and build predictive customer models. Merchants lose actionable insights: they don't know why customers chose competitors, which features influenced decisions, or how agents ranked their products against alternatives. While agents streamline shopping, they threaten traditional retail analytics, forcing merchants to adapt or face competitive disadvantages in a data-scarce ecosystem.

Merchant Impacts

Merchants used to leverage direct interactions for data on user journeys, enabling targeted upselling and risk assessment, but AI agents now display only tokenised transactions without context like end user behaviour or preferences. This blindness prevents merchants from:

1) Detecting fraudulent agents or synthetic transactions, as they cannot distinguish legitimate shopping patterns from malicious bots;
2) Optimising product listings, as they don't know which features agents prioritised;
3) Building retention strategies, as they have no customer preference data;
4) Conducting cohort analysis or customer testing, as agent-mediated purchases appear as anonymous transactions.

Practical Applications

In fashion e-commerce, a lack of behavioural insights prevents merchants from analysing why agents select certain items, missing opportunities to refine sizing recommendations or detect return fraud.

In travel booking, opaque agent decisions obscure user intent (e.g., budget vs. luxury), hindering personalised upsells and exposing to disputes from misunderstood bookings.

A merchant can no longer observe which products agents compare (competitor A vs. competitor B), which product attributes agents prioritise (price vs. durability vs. sustainability), or where agents abandon comparisons (e.g., when merchant's product is 15% more expensive than a competitor). This intelligence gap means the merchant cannot refine product positioning, adjust pricing strategy, or highlight differentiating features that agents might value.

Call to Action

Some platforms may share anonymised behavioural insights in exchange for merchant priority. Merchants should negotiate with these platforms to gain first-party data access, which they can use to position their product accordingly, or detect fraudulent agents. Merchants who begin building alternative data streams now will maintain customer intelligence; those who delay will lose visibility permanently, becoming unable to compete on personalisation or retention.

Click to return to table of contents
Click to return SWOT summary
Click to return summary of threats

Reduced Impulse Purchases – an agentic commerce threat

Summary

In agentic commerce, reduced impulse purchases stem from AI agents' goal-driven efficiency, which compresses the customer journey by automating targeted buys without browsing, emotional triggers, or lucky discoveries, reducing impulse-driven revenue especially in categories like fashion or snacks. Merchants face eroded margins as traditional tactics like flash sales lose effectiveness, shifting monetisation to commissions on agent platforms. This hyper-rational agent behaviour strips away the psychological triggers that fuel impulse buying: perceived urgency (limited-time offers, scarcity messaging), emotional engagement (aspirational imagery), and irrelevance discounting (agents ignore products outside user preferences). Though agents enhance planned shopping, they disrupt impulse models, and collapse margin structures built on emotional purchases – rendering traditional promotional mechanics ineffective.

Merchant Impacts

Traditionally, merchants drove margin through impulse purchases, where retailers used psychological triggers (scarcity, limited-time offers, FOMO) to accelerate this behaviour. In agentic commerce, this revenue stream disappears. Agents are emotionally indifferent; they do not experience FOMO, do not respond to aspirational imagery, and do not impulse-add items to cart. Instead, agents systematically eliminate products outside user-defined criteria, refuse add-ons that don't improve value, and exploit promotional mechanics aggressively (applying every available discount, stacking loyalty points, price-comparing in real-time). High-margin categories (apparel, accessories) face revenue losses as impulse volumes fall dramatically.

Practical Applications

A grocery delivery service historically captured revenue from impulse add-ons (customers buying ice cream, crisps, sweets on top of planned purchases). When shopping through an agent, customers specify a list of products at a certain total cost. The agent orders those exact items, within the price limit, and then stop, adding no impulse items. Across the customer base, add-on revenue (which was pure margin) collapses.

In electronics, agents compare prices cross-platform without emotional prompts, bypassing flash deals that trigger impulse gadget buys.

Call to Action

It is difficult to combat this threat since the whole purpose of AI agents is to streamline the customer purchase journey. One approach could be for merchants to implement dynamic pricing APIs that allow agents to apply time-sensitive offers, drawing from real-time inventory data to create scarcity alerts that mimic traditional impulse triggers. Merchants can also develop their own, retailer-owned AI assistants, that integrate into platforms like chat apps for proactive queries that fill the same role as impulse purchases for the merchant. Retailers should invest promptly in these adaptations to revive impulse mechanics, training models on behavioural data to bridge gaps and prevent rivals from dominating with their own agent systems.

Click to go SWOT summary

Click to return to table of contents

09

Edgar, Dunn & Company

Point of contact

edgar, dunn & COMPANY

Edgar, Dunn & Company (EDC)

Independent Global Financial Services and Payments Strategy Consultancy

Founded in 1978, the firm is widely regarded as a trusted advisor to its clients, providing a full range of strategy consulting services, expertise and market insight, and M&A support

EDC has been providing advisory services to the whole payment value chain, including …

Top technology powerhouses, social media platforms, and marketplace facilitators

Consumer and business payments and financial services

All major international card associations / schemes & many domestic card schemes

Many of the world's leading retailers, travel merchants, and major airlines

More than 40 European banks, lenders, Fintechs, payment companies

Most of the top 25 financial institutions and Fintechs

Many of the world's most influential payments providers and processors

Edgar, Dunn & Company (EDC)

Independent Global Financial Services and Payments Strategy Consultancy

Office Locations

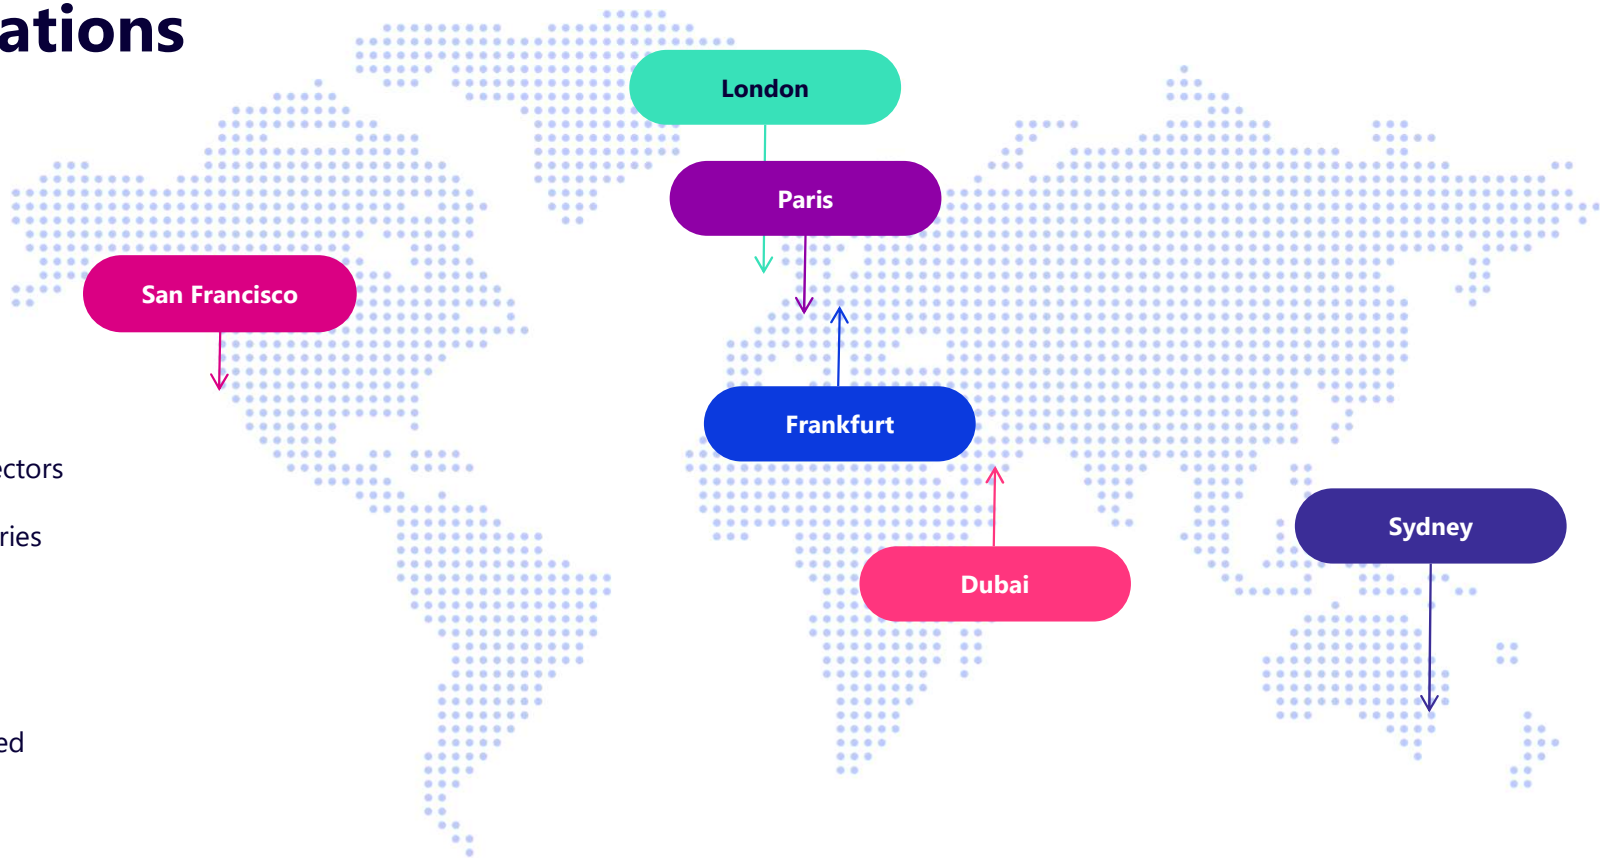

- **Fintech & Payment Strategy & M&A**
- **Independent – Owned** & Controlled by EDC directors
- **+250 clients** in 55 countries and 6 continents
- **6 Office** Locations Worldwide
- **+1800** Projects Completed

edgar, dunn & COMPANY

Point of contact

Retail & Hospitality Practice

Mark Beresford
mark.beresford@edgardunn.com

Global Sales & Marketing Manager

Amine Saidi
amine.saidi@edgardunn.com