



Exploring the AI-powered commerce shift underway

– and how to compete
when AI guides
the purchase

AI attention, human loyalty, or both?

AI is quickly changing how we live, work, and play.

For many, to live means to shop. And the new way to shop more often involves AI choosing the route after customers set the destination.

AI can participate in shopping decisions at different levels. In some cases, it acts as a research assistant that gathers information and summarizes options. In others, it narrows a large set of possibilities into a shortlist of recommendations. Increasingly, it may also complete transactions on the consumer's behalf when given permission to do so. These are distinct stages of involvement, and consumer

willingness to delegate varies across them.

In 2025, already one-quarter of consumers used generative AI shopping tools, and agentic AI is moving into routine tasks. For example, a little over half of consumers use virtual assistants to automate re-ordering or meal planning at least once a week, according to a Capgemini Research Institute survey.¹

While brands compete for a top spot in AI recommendations, these companies (especially

1/4

of consumers used generative AI shopping tools in 2025

direct-to-consumer players) need to earn an equally prized place in the hearts of customers.


Because who follows their GPS every time? Sometimes people want a detour, or the familiar road, scenic route, or dusty trail that adds an unexpected stop no algorithm would have chosen. Shopping decisions are beginning to work in much the same way.

Consider Jack, a typical consumer in search of the perfect pair of running shoes that match his size, preference, and budget. He prompts his AI agent: "Find me a quality pair

of dark blue running shoes under \$130 in a size 11. And make sure the fit is a little wider than normal and suitable for casual wear."

With no particular brand specified, his agent simply returns a list of shoes matching the request, scraped from well-structured websites, of course. Some brands are there, while others are not. Those not appearing in the list might assume all is lost. But not so fast. If Jack asks, "What about Brand X?" he forces a U-turn and a refresh of the list. Suddenly, the conscious override keeps a favorite brand in the running.





Earn preference inside the agent's logic and inside the customer's mind

There's a broad shift in how commerce is being designed and experienced.

For years it's been built around a familiar model. Brands and retailers designed journeys, attracted traffic, and guided customers through search and comparison in hopes their attention led to a purchase.

To win in the new agentic commerce landscape, organizations now have to solve a two-pronged challenge: "How can we become visible, understandable, and transactable for AI agents

while still building brand appeal and loyalty with the people those agents serve?"

An audience of one, the human consumer, was yesterday's reality. Today, agentic AI is becoming a new participant in the shopping journey, demanding attention in its own right. As AI agents increasingly mediate interactions, from discovery to purchase, brands are gradually losing their grip on direct engagement with consumers.

And today's discovery-to-purchase use cases are only the first chapter.

The next wave of AI will move beyond recommendations and transactions to anticipating consumer needs, making decisions, and orchestrating purchases before customers even realize they need them.

Capturing this opportunity goes beyond just ranking high in search results, especially when traditional methods such as SEO (search engine optimization) can no longer guarantee visibility. Instead, brands need GEO (generative engine optimization), a new discipline for structuring content and data so AI systems can read, trust, and recommend them.

Sites that are not machine-readable or structured for AI risk becoming invisible from what the AI recommends. In contrast, organizations that

make their content easily interpretable by AI are gaining disproportionate visibility and influence.

Adobe, a leader in creative technology, has found that traffic from AI sources to US retail sites grew almost 400% year-on-year in Q1 2026, following an almost 700% increase during the 2025 holiday season. More importantly, AI-referred visitors are proving far more valuable: they convert 42% better than traditional traffic sources and exhibit higher engagement and longer time on site.²

But earning a place in AI recommendations requires more than being readable to machines. Brands and retailers must understand the consumer problems, occasions, and decision criteria that agents are likely to evaluate on behalf of shoppers. Those insights then need

to be reflected consistently across products, packaging, pricing, fulfillment options, and supporting content. The brands most likely to earn recommendation are those that align what consumers value with the signals agents use to judge relevance and suitability.

Traffic from AI sources to US retail sites grew almost

400%

year-on-year in Q1 2026.

AI-referred visitors convert

42%

better than traditional traffic sources.

² Adobe report: US retailers see surge in AI traffic, but many websites aren't entirely readable by machines

The agentic commerce wake-up call for consumer products and retail (CPR)

According to some industry estimates, the agentic AI market in retail and e-commerce could grow from about \$60.4 billion in 2026 to \$218.4 billion by 2031.⁽³⁾ This reveals a simple truth; If your products aren't fully visible or accessible to AI agents, your brand revenue may be at risk.

CPR is a natural starting point for agentic commerce, but not every category will be affected in the same way. The sector sits on a sliding scale of consumer preference. At one end are high-frequency repeat purchases such

as groceries, household goods, and personal care items. Here, consumers often seek the best combination of price, availability, and convenience rather than actively evaluating every purchase. These categories are ideal candidates for delegation to AI agents. Why? Well, brand appeal is not the decisive buying factor for consumers, so these brands must instead redirect more effort to appealing to AI agents.

At the other end are prestige beauty, luxury, and identity-driven categories. Here human preference

remains strong, and consumers are more likely to override an agent's recommendation for a brand they know and desire.

Recognizing this split matters because CPR organizations don't need one averaged response to agentic commerce. They need category-specific strategies that favor machine selection where loyalty is thin and human preference where brand meaning still carries decisive weight.

³ Mordor Intelligence: Agentic AI in retail and e-commerce market size & share analysis



The agentic AI market in retail and e-commerce could grow from about

\$60.4 B in 2026

to

\$218.4 B by 2031



Not every conversion will be won the same way

Four distinct shopping occasions reveal how much control consumers may be willing to hand over to an agent, and when brand preference may pull them back. Override behavior will vary across these occasions depending on the consumer, purchase mission, channel, or context.

1 Routine choice

High delegation,
low brand attachment

Routine replenishment where price, availability, and convenience matter most. Agents often decide outright, making machine visibility, availability, and pricing the primary competitive levers.

2 Default preference

High delegation,
strong brand attachment

Established favorites such as coffee, detergent, or pet food. Consumers delegate the purchase but expect the agent to preserve their preferred brand, making retention and machine-readable brand superiority critical.

3 Discovery

Low delegation,
low brand attachment

First-time purchases, gifts, and exploratory buying. Agents act as advisors, shaping preference through evidence, reviews, comparisons, and product attributes.

4 Identity and passion

Low delegation,
strong brand attachment

Prestige beauty, enthusiast products, and other identity-driven purchases. Consumers remain closely involved, and strong brand equity can motivate overrides of an agent's recommendation.

Importantly, these behaviors do not map neatly to companies or even brands. A broad portfolio can occupy several positions at once, and the same product can shift depending on the shopper, occasion, channel, or task. The strategic unit is therefore not the organization as a whole, but the specific combination of brand, product, shopper, occasion, and purchase mission.

Different players, different levers

This distinction resolves a key paradox. The categories most exposed to automation are often the ones where overrides are least likely.

As a result, an override-led strategy applies to only part of the portfolio. Elsewhere, success depends on defaults, evidence, availability, and machine legibility.

Sellers need to quickly acknowledge that customer expectations are changing. Conversational interaction is becoming the default as quick AI queries gradually replace endless comparison browsing. So, how can organizations respond?

Player	Controls	Opportunity	Risk
Retail	Interface, customer relationship, first-party data	Become the trusted environment where preferences are captured and reused	Becoming fulfillment infrastructure for external agents
Consumer-products manufacturers	Products, brands, content, demand generation	Build machine visibility and create preference before recommendations are made	Losing influence inside agent experiences they do not control
Marketplaces	Discovery, catalog access, placement policies	Shape visibility through rules, protocols, and emerging agent commerce models	Erosion of trust if commercial incentives outweigh relevance
Direct-to-consumer brands	Owned channels, customer data, direct relationships	Create agent-ready experiences while observing customer behavior directly	Discovery increasingly starts outside owned channels

Be seen, be understood, be chosen (twice!)

Many are already extending their existing commerce models and building capabilities to operate across two interconnected decision layers: how agents evaluate options, and how humans experience and validate those decisions.

The three-point checklist below is essential for earning visibility with agents and maintaining appeal with humans. While being visible and understandable are the foundations of GEO, being

transactable is what allows sellers to take advantage of agentic commerce.

1

Be seen – ensure agent visibility

Being visible is no longer guaranteed. Product content that's visually stunning to humans but technically unattractive to machines may not be discoverable and readable by AI systems.

Diagnostic tools like AI Content Visibility Checker from Adobe can help analyze how readable a website is by machines.

2

Be understood – provide reliable context

Agents operate on intent, not keywords. They need reliable, real-time context to make decisions. That requires unified product, pricing, inventory, and customer data as well as a semantic layer to give agents a consistent basis for interpretation.

The semantic layer can be thought of as a translator sitting between agents and enterprise systems. Every

source system may use different labels, structures, and terminology, but the semantic layer translates them into a shared business language. As a result, when an agent encounters terms like “customer,” “account,” “offer,” or “availability,” it understands their meaning, regardless of where the data originated.

3

Be chosen (by the machine and by the human) – enable agent transactions and human override

Before an AI agent recommends a product, it may evaluate signals from hundreds or even thousands of digital sources, including retailer sites, marketplaces, reviews, social communities, publisher content, product databases, comparison engines, and brand-owned channels.

In that context, a brand's website may represent only a tiny fraction of the evidence shaping the recommendation. Being chosen therefore depends not only on what a company says about its own products but also on how consistently and credibly it is represented across the broader digital ecosystem. The battle for visibility is shifting from controlling a

destination to influencing the network of sources that AI uses to determine relevance, trust, and value.

Beyond these product "quality" signals, selection will also increasingly depend on whether a product can move from recommendation to transaction inside an agent environment. That requires more than making product content visible and understandable. Brands and retailers will need real-time APIs for functions like pricing, availability, fulfillment, checkout, and payment, supported by protocols that allow internal and external agents to communicate safely and consistently.

Ultimately, being chosen is not only a machine decision. Consumers still get to decide when to delegate, when to trust, and when to override an agent's answer. Even though AI can bring up the cheapest, fastest, or most logical option in seconds, people can easily shoot it down with a "no, I want that one instead." That's why brands need to preserve the emotional, experiential, and identity-based pull that drives customers to willingly defy AI (at least some of the time).

Is there a standout transaction protocol?

Not yet. The protocol layer is still being contested, with different ecosystems pushing different routes into agent-enabled commerce. Whether it's Google and Shopify's Universal Commerce Protocol (UCP), OpenAI and Stripe's Agentic Commerce Protocol (ACP), or another one that comes out on top, the practical thing to do now is avoid betting too early on a single winner. Instead, build a protocol-agnostic commerce foundation that's flexible enough to adapt as the standard-setting race continues.

Two dimensions to make agentic commerce real

The first is a question of technology (the seemingly easy part).

Every site needs an agentic door to enable communication with AI. But before that can happen, organizations need to make their commerce stack agent-ready by resolving fragmentation across ERP, CRM, supply chain, and content systems. Then as they build the semantic layer, they can provide

API-driven access to product, price, inventory, and fulfillment data.

This foundation should be complemented by orchestration, tool registries, memory, monitoring, and control mechanisms that allow agents to act safely without replacing existing systems outright.




The second and harder dimension is team transformation.

Agentic commerce blurs the lines between marketing, sales, technology, data, risk, legal, and even HR as organizations move from human-only workflows to hybrid models where people supervise networks of agents.

Roles that once had clear separation will start to collide. For example, the CMO can no longer treat brand, content, and campaigns as separate from APIs, data structures, and product intelligence. Conversely, the CIO can no longer

treat architecture, uptime, and governance as purely back-office concerns.

Today the API is the marketing campaign, and the data structure is the brand story. The organizations that win this will either merge the two roles into one accountable owner, or build a new seat explicitly chartered to hold both. A brand strategy and data architecture that aren't designed by the same function will keep producing a disjointed experience.



Measure the recommendation – and the override

Three areas will matter most when measuring performance within agentic commerce:

Track agent visibility

How often do your products appear in AI-generated recommendations?

Assess agent representation

How are your products described, compared, and ranked?

Monitor agent-driven performance

What's the conversion rate, basket size, and margin impact?

Capgemini helps organizations use simulation environments and synthetic users to continuously test how agents perceive and recommend their products. These simulations show whether a brand is visible, understandable, and selectable to agents, and where its content, data, and architecture need to improve.

But agent performance is only half the story. What about when consumers reject the agent's recommendation in favor of a "detour" – a specific, named alternative?



This is where conscious override rate (COR) becomes useful as a measurement concept.

COR may be measured directly, estimated through proxy signals, or inferred through consumer research, depending on channel and agent architecture.

But COR isn't a typical live dashboard metric. Unless the override happens within a retailer's own agent, brand app, or site, it's invisible and cannot be measured as a clean event stream. Inside third-party agents, sellers typically see the order but not the deliberation. But that doesn't mean there aren't signals to follow.

The starting point is a stated-preference survey instrument that estimates how often consumers would pass over an agent's recommendation for a named brand. Between survey waves, organizations can also monitor proxy signals that suggest whether override behavior is rising or

falling. Examples include branded query share after agent exposure (are more consumers searching for your brand by name after receiving an AI recommendation?), direct conversion following a third-party recommendation, and the share of prompts carrying a brand constraint, such as "only show me Brand X." These signals can be captured through partner instrumentation.

But this kind of measurement is still not enough. We must also look for social signals of overrides. If Jack quietly ignores his agent's recommendation and buys the shoes he wanted all along, the brand wins the sale, but nobody else learns from it.

Now, what if Jack tells a friend why he skipped the "logical" choice, leaves a review, shares an

unboxing, or posts about the product? Details of the detour now begin to emerge. And it's also where human loyalty and AI visibility start to reinforce each other. The content consumers create when they explain their choice is also content that AI systems can see and use.

One hypothesis worth testing is that authentic overrides may strengthen the evidence base available to future agents. When consumers choose a particular brand, they create signals that AI systems may subsequently use when evaluating recommendations. If this relationship proves consistent over time, the loyalty flywheel and the visibility flywheel may become mutually reinforcing.



Poetry for humans, evidence for machines

A luxury beauty brand discovered that strong market positioning didn't automatically translate into AI visibility.



Through a large-scale GEO analysis, Capgemini simulated more than 12,000 interactions across multiple platforms (Gemini, ChatGPT, Perplexity), customer personas, and purchase journeys. Our findings revealed a striking contrast: while the brand was frequently recommended in fragrance, it was largely absent from skincare and makeup conversations.

The root cause wasn't the product quality or brand awareness. It was how the brand's expertise was represented. Since AI systems favor structured, evidence-based content, the emotional and sensory

language around the brand's heritage, craftsmanship, and product performance was largely overlooked.

Capgemini worked with the client to identify and prioritize a set of actions focused on translating existing strengths into machine-readable signals, including structured product data, clinical evidence, FAQs, and comparison content.

Forward-looking goals include a 67% increase in fragrance share of mention among specialist consumers, a doubling of visibility in selected makeup journeys, and a move from near-zero visibility to meaningful presence in skincare recommendations.

The project revealed an emerging reality of agentic commerce: organizations must not only communicate effectively with people but also ensure that AI systems can understand, interpret, and confidently recommend them.

Forward-looking goals include a

67%

increase in fragrance share of mention among specialist consumers

Cheat sheet

Five decisions for competing in agentic commerce.

Together, the shopping occasions and player roles show which purchases are most likely to be delegated, where human preference still matters, and which levers each type of company can control. The next step is to turn that diagnosis into five connected decisions:

Choose where to compete

Identify which occasions must win the default, the recommendation, or the override.

Act within two quarters

Focus on one or two priority categories, establish baselines, fix the data foundation end-to-end, and prepare for emerging commercial models.

Build the enabling capabilities


Strengthen all three links in the preference chain through preference capture, machine-readable content, semantic data layers, APIs, and agent orchestration.

Assign ownership

Align marketing and technology around shared accountability for how the company is represented and performs inside agent ecosystems.

Decide how you will know

Track agent visibility, representation, performance, override behavior, default retention, and named demand.



The route is changing. The destination isn't



The route is changing. The destination isn't.

Agentic commerce isn't an incremental evolution of e-commerce. It fundamentally changes how purchasing decisions are made.

Customers still set the destination. They still define what they need, what they value, and what they trust. What's changing is how they get there. AI may help gather information, narrow options, recommend products, or complete transactions on their behalf. What changes from purchase to purchase is how much of that journey consumers choose to hand over to the agent.

For brands and retailers, the challenge is becoming visible and understandable to the agent without becoming forgettable to the customer.

The winners won't be those that optimize solely for algorithms or rely solely on brand equity. They'll be the organizations that earn a place in the recommendation and remain desirable enough to survive the detour.

Capgemini helps organizations prepare for that future.

Want to learn how to compete for both machine recommendation and human preference? Contact us today.

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real | www.capgemini.com

Authors



Dreen Yang

EVP, Global Consumer Products and Retail Lead,
Capgemini
dreen.yang@capgemini.com



Maxime Girardeau

VP, Global Commerce and GEO Offer Lead,
Capgemini
maxime.girardeau@capgemini.com



Mark Ruston

VP, Global Retail Lead,
Capgemini
mark.ruston@capgemini.com



Sebastien Jouvenaar

VP, Strategy and Innovation, Head of the Applied
Innovation Exchange, frog - Part of Capgemini Invent
sebastien.jouvenaar@frog.co

