QuantumBlack
AI by McKinsey

The automation curve in agentic commerce

Agentic AI is increasingly a part of shopping, but not all transactions will be automated in the same way. Here's what agents will handle—and the situations that will call for human involvement.

This article is a collaborative effort by Deepa Mahajan, Hannah Mayer, Katharina Schumacher, and Roger Roberts, with Katharina Giebel, representing views from QuantumBlack, AI by McKinsey.

January 2026

For many shoppers, this past holiday season may have felt different. Perhaps an AI assistant suggested gifts your relatives might actually like while filtering for items that could arrive before the holidays. Maybe it helped you navigate the specs of three different noise-canceling headphones or scanned five retailers for a specific holiday outfit, assembled a ready-to-buy basket, and politely asked, "Should I go ahead?"

This is the year AI agents stopped being an experiment and became part of how people shop, not in headline-grabbing ways but in everyday moments—helping shoppers make sense of choices, assemble baskets, resolve trade-offs, and move toward action. Yet what looks like small convenience today is an early signal of a much larger shift in the way we shop. According to our research, even under moderate scenarios, AI agents could mediate $3 trillion to $5 trillion of global consumer commerce by 2030.[1] Because agents navigate the same internet as humans—visiting websites, engaging with APIs, and interacting with loyalty programs—they can scale quickly. And as they do, they are reshaping how intent forms, how products are discovered, and where value pools can be found.

We introduced many of these themes in our report *The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants* last fall. This article builds on that foundation. Here, we explore what we call the "agentic commerce automation curve," which illustrates how the shopper experience shifts at different levels of delegation, and outline how retailers can prepare for a world in which the customer is still human but AI agents increasingly mediate key decisions.

The six-level agentic commerce automation curve

The rise of agentic commerce reflects the collision of three forces. First, AI agents have reached decision-grade usefulness, allowing consumers to delegate not only inspiration but also shortlisting, assembly, and even execution. Second, the ecosystem now has rails for real autonomy. Open-source protocols—such as MCP, A2A, AP2, ACP, and UCP—enable agents to read data, negotiate with other agents, and transact safely. The Linux Foundation recently established the Agentic AI Foundation—a partner-backed effort including Anthropic, Block, Google, Microsoft, OpenAI, and others—focused on the interoperability, identity, and payments building blocks needed to make autonomous commerce viable at scale.[2] Third, intent is shifting upstream. Agents increasingly act when consumer goals surface—such as a conversation about an upcoming birthday party, a calendar reminder for a trip, or a low-supplies signal from a device. For retailers, the implications are stark: If your catalog, policies, and value proposition are not machine-readable, agents—and by extension, shoppers—simply will not find you, no matter how beloved your brand is.

That said, the rise of AI agents does not represent a single leap from human-driven shopping to full autonomy. Instead, agentic commerce is unfolding along a curve—one defined by how much of the commerce journey consumers are willing to delegate to machines. This automation curve is composed of six distinct levels of automation, each representing a different mode of

[1] *The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants*, McKinsey, October 17, 2025.
[2] "Linux Foundation announces the formation of the Agentic AI Foundation (AAIF), anchored by new project contributions including Model Context Protocol (MCP), goose and AGENTS.md," Linux Foundation, December 9, 2025.

delegation—from basic rules-based convenience to fully autonomous multiagent coordination. Importantly, these levels describe what agents are technically capable of doing, not what consumers will always choose to allow.

How agentic commerce plays out in B2B

At first glance, B2B commerce may seem fundamentally different from consumer shopping. Corporate buying decisions, after all, often involve committees, contracts, compliance, and months-long cycles rather than impulse and inspiration. But when viewed through the lens of agentic commerce, the same automation curve applies. It's the stakes, constraints, and optimization targets that change.

In consumer commerce, delegation is personal. A shopper authorizes an agent to save time or money. In B2B, delegation is institutional. Authority flows from procurement policies, budget owners, risk teams, and legal frameworks. As a result, autonomy advances more slowly. But once unlocked, it scales far more powerfully.

At early levels, agents assist buyers by synthesizing complex information and by comparing suppliers against technical specifications, historical performance, sustainability criteria, and contract terms. Assembling a "basket" means constructing a procurement-ready package: products, volumes, pricing tiers, service levels, and obligations ready for approval rather than checkout.

The real inflection point comes at supervised execution. When agents are authorized to act within clearly defined policies—such as spending thresholds, preferred vendors, and compliance rules—they can manage replenishment, renewals, substitutions, and exceptions automatically, escalating only when something breaks. Here, autonomy optimizes reliability and continuity, not convenience.

At higher levels, agents can operate on standing enterprise goals codified by leadership, audited by policy, and bounded by reversible authority. Examples include maintaining uptime at the lowest risk-adjusted cost, balancing supplier concentration, or hitting environmental, social, and governance targets without disrupting operations. Competition shifts from unit price to predictability, transparency, and policy alignment.

In consumer commerce, autonomy feels like magic. In B2B, it feels like governance finally catching up with complexity and turning it into an advantage. For enterprise leaders, the question is no longer whether AI can automate buying but whether governance systems are ready to delegate authority without losing control.

Further, adoption will not necessarily move uniformly "up" the curve. While agentic capabilities continue to advance, two forces are shaping consumer delegation. The first is time and trust: As consumers gain familiarity with agents and see them perform reliably, they become comfortable delegating larger portions of the journey. The second is category dynamics. Willingness to delegate varies sharply by ticket size, emotional salience, identity signaling, and regret risk. (This article focuses on the retail experience; for a look at how agents could impact B2B commerce, see sidebar, "How agentic commerce plays out in B2B.")

Together, these forces determine a ceiling of delegation, in which autonomy naturally plateaus for a given category or moment. In some contexts, consumers may be comfortable delegating end-to-end execution. In others, they will deliberately stop short, retaining control not because agents are incapable but because human involvement is intrinsic to the value of the experience.

For these reasons, the model is best understood as a curve rather than a ladder. Higher levels of automation are not inherently better or more advanced, and the goal is not maximum autonomy but optimal delegation (exhibit).

Level 0: Programmed convenience ('set it and forget it')
This level is the pre-agentic baseline: Recurring replenishment for things that run out—coffee pods, detergent, diapers, shampoo—is handled through subscriptions, scheduled refills, and recurring shipments. At this point on the curve, automation is rules-based—useful but brittle and largely blind to context. When needs change, it breaks, and the human steps back in.

Still, level 0 proves a foundational point at which consumers delegate when automation is reliable and reversible. For example, around 23 percent of US Amazon shoppers had at least one active Subscribe & Save order in 2024.[3]

Level 1: Assist ('the cognitive sidekick')
At level 1, agents help shoppers think and make decisions, but they do not execute. A shopper might ask, "Find four gifts under $75 that can ship by Friday; prefer sustainable brands; and summarize trade-offs." Or, in a more complex category, "Compare three noise-canceling headphones, and explain how they differ on sound quality, battery life, and comfort." The agent's role is analytical. It scans catalogs, parses reviews, compares features, and synthesizes options into short lists or recommendations. Crucially, it does not commit to a configuration or resolve operational constraints. There is no cart, no basket, and no readiness to transact. The human evaluates the options, weighs trade-offs, and decides what to do next.

In other words, level 1 replaces search and comparison but leaves assembly and execution entirely with the shopper.

Implications for retailers: Verifiable data beats marketing gloss. Agents require information they can parse and compare—structured attributes, clear eligibility rules, sizing and fit certainty, and claims that can be substantiated.

[3] *CIRP - Amazon Report,* "A surprising percentage of Amazon customers have or had a Subscribe & Save recurring order!," blog entry by Michael Levin and Josh Lowitz, Substack, February 18, 2025.

Exhibit

Automation and delegation in commerce will not unfold uniformly, and they will vary by category and required human involvement.

LOW AUTONOMY ↑ HIGH AUTONOMY

Automation level	Agent role	Human control	What it looks like	Example transaction
0 **Programmed convenience** (pre-agentic baseline)	N/A (rules-based subscriptions)	Consumer sets all guidelines for transactions	Prescheduled refills or rides set automatically (eg, auto-ship)	Automatically delivers 1 pound of coffee beans every 2 weeks
1 **Assist** (cognitive sidekick)	Gathers and presents information	Agent offers advice only; consumer completes transactions	Agent provides product summaries or comparisons	Agent compares 2 noise-canceling headphones, summarizing trade-offs on sound quality, battery life, and comfort
2 **Assemble** (personal shopper)	Assembles orders and stops at "ready to buy"	Consumer receives agent-assembled baskets and can approve for purchase	Agent assembles ready-to-buy baskets of goods or services and presents alternatives, ETAs, etc	Agent builds a ready-to-book weekend trip including flights, hotels, and activities and presents for approval
3 **Authorize** (supervised executor)	Executes orders within consumer-defined guardrails	Consumer authorizes agent to complete purchase based on explicit instructions	Agent automatically places order within guardrails (eg, price, ETA), providing alternatives if conditions cannot be met	If preferred groceries cost less than $200 and can be delivered Friday, agent places the order; otherwise, agent asks consumer
4 **Autonomize** (intent steward)	Operates on standing intents, anticipating, negotiating, and resolving issues	Consumer sets mandate and schedule (eg, budget, scope, expiration), with power to audit or change	Agent operates on standing intents, anticipates needs, and autonomously replenishes or services within defined goals	Agent maintains airline loyalty status at the lowest possible cost over the year, rebooking flights and handling changes automatically
5 **Networked autonomy** (multiagent commerce)	Negotiates and coordinates autonomously with other agents	Minimal ongoing human intervention; control exercised through predefined mandates and auditability	Agents coordinate across platforms to source, negotiate, and fulfill services with minimal human input	Personal agent negotiates broadband, energy, and insurance contracts across providers, making changes within predefined mandates

McKinsey & Company

Level 2: Assemble ('the personal shopper')

Level 2 marks a qualitative shift: Agents move from analysis to orchestration. Here, the shopper expresses an intent, and the agent returns a purchase-ready basket. "Build a cozy winter outfit

under $150." "Stock a pantry for a vegan guest arriving tomorrow and staying for three days." Or, more complex, "Put together a home office setup under $2,000 that supports dual 4K monitors and quiet video calls and has next-day delivery." Unlike level 1, the agent is tasked with resolving trade-offs and constraints rather than merely surfacing them. It selects specific items, ensures technical compatibility, and balances performance against price, availability against delivery speed, and promotions against eligibility. Taxes, shipping windows, loyalty benefits, and substitutions are handled by default. The output is not a list of options; it is a coherent configuration that is ready to check out. The shopper's role shifts accordingly from comparing options to approving or adjusting a proposed solution.

Implications for retailers: Success at level 2 requires API-first merchandising. Inventory, pricing, shipping promises, promotions, and returns logic must be exposed cleanly so agents can assemble baskets with human-level fidelity.

Level 3: Authorize ('the supervised executor')

At level 3, consumers delegate not only actions but also rules. Instead of approving each step, they authorize an agent to execute within clear boundaries. "If groceries are under $120 and arrive Friday 6–8 p.m., place the order." "If my preferred sneakers drop below $80 from merchants I trust, buy them." The agent then runs the workflow end to end, choosing among eligible options, swapping out-of-stock items for approved substitutes, applying loyalty benefits, and escalating to the shopper for approval only when something falls outside the rules.

Implications for retailers: To support shoppers at level 3, merchants must make it possible for an agent to pay and act on a customer's behalf with safety and transparency. That means purchasing authorization that can be limited (by budget, time window, merchant, or category), activity that can be audited (what was bought and why), and actions that can be reversed (easy cancellations, refunds, and overrides when needed).

Level 4: Autonomize ('the intent steward')

At level 4, agents operate against standing goals rather than one-off transactions. For example, "Keep household essentials under $300 per month." "Maintain my airline loyalty status at the lowest total cost over the course of 2026." "Make sure we never run out of baby supplies." The agent continuously monitors needs, anticipates replenishment, compares options across merchants, and optimizes for longer-term outcomes such as maintaining or achieving a certain loyalty status. The agent then handles the operational follow-through, including changes, returns, and replacements. The shopper becomes episodic, stepping in mainly for meaningful decisions or exceptions.

Implications for retailers: Competition at level 4 shifts from winning a single purchase to earning a place in the agent's ongoing plan. Merchants need deeper integration—especially around loyalty, eligibility, substitutions, and service guarantees—so agents can reason about trade-offs and execute reliably. Put simply, it's no longer enough to expose a catalog; retailers must expose the rules and policies that determine what "good" looks like.

Level 5: Networked autonomy ('multiagent commerce')
This forward-looking level is still emerging and points to a world in which commerce becomes agent-to-agent by default. Personal agents won't just interact with merchant websites; they will negotiate directly with a network of specialized agents that optimize pricing, logistics and delivery, payment authorization, and loyalty programs. Ultimately, this will result in multiagent marketplaces where intent can be brokered, trust is carried through reputation signals, and transactions are settled through shared protocols—enabling "procurement as a service" to run continuously in the background.

Implications for retailers: Level 5 will be shaped by those that are already proficient at level 4. Retailers that expose policies, guarantees, and loyalty logic in machine-readable ways will be positioned to influence how these ecosystems route demand. Those that don't risk becoming interchangeable suppliers competing primarily on price in machine-negotiated flows.

How the automation curve bends: Where delegation accelerates, plateaus, and reshapes value

The automation curve describes what AI agents can do across the shopping experience. It also can help explain the way that delegation can play out in practice and why automation does not unfold evenly across categories, moments, or consumers.

In the real world, consumers do not climb the curve uniformly, nor do they aspire to full autonomy in every context across shopping categories. Instead, delegation accelerates where automation removes friction without sacrificing meaning. It plateaus where human involvement is intrinsic to value, and it becomes selective amid trade-offs and uncertainty. Understanding these patterns is critical for retailers deciding where to invest, what to expose to agents, and how to compete in an agent-mediated world.

Where delegation accelerates: Utility, repetition, and low-regret purchases
In categories where shopping is primarily a task rather than an experience, delegation tends to move quickly up the curve. Low-regret purchases such as groceries, household essentials, and basic consumables are natural candidates for higher autonomy. Here, the value of shopping lies in efficiency, reliability, and predictability rather than discovery or expression.

As agents prove capable of assembling baskets accurately, executing within guardrails, and handling substitutions or delivery changes gracefully, consumers become comfortable delegating execution entirely. Attention shifts from evaluating options to reviewing outcomes: Was the order on time? Did it stay within budget? Were substitutions reasonable? Over time, approval becomes implicit and intervention becomes the exception.

For retailers, this dynamic reshapes competition. Brand storytelling and front-end experience matter less than operational trust. Agents optimize for delivered value—factors such as price, availability, service reliability, and reversibility. Merchants that expose clean inventory data,

predictable fulfillment performance, and transparent substitution and return policies become default suppliers, often without ever "winning" a traditional moment of consideration. In these categories, being agent-readable and dependable matters more than being distinctive.

Where delegation plateaus: Identity, aspiration, and regret risk

In high-consideration categories, such as luxury goods or milestone purchases, delegation often plateaus lower on the curve. Here, shopping is not merely about outcomes; it is about identity, intent, and emotional assurance. Consumers may enthusiastically enlist agents to research, compare, and analyze but stop short of fully autonomous execution.

Consider a luxury handbag purchase: A consumer may ask an agent to evaluate how different brands hold value over time, analyze resale markets, or assess how a particular style aligns with their personal aesthetic. The agent may surface alternatives, identify better price points in the resale market, or locate in-store availability. But the final decision and the transaction itself remain firmly human.

In these moments, the agent functions less as an executor and more as an analyst and curator. The ceiling of delegation is set not by technical limitations but by emotional and identity-based considerations, such as the desire for a tactile experience, social signaling, or the avoidance of regret. Importantly, lower autonomy does not imply lower value. In many such categories, human involvement is itself a key component of the product.

For brands, this distinction is critical. Competing effectively does not require pushing consumers toward full automation. It requires enabling agents to support deliberation by exposing rich contextual attributes, provenance, craftsmanship, and long-term value signals while preserving human control at the point of commitment. In these categories, winning means shaping how decisions are informed, not how quickly they are executed.

Where delegation is selective: Complexity, trade-offs, and context

Most categories sit between these two poles. In travel, consumer electronics, home goods, and other complex purchases, delegation is selective and situational. Agents may autonomously handle research, comparison, monitoring, and assembly while escalating decisions that involve meaningful trade-offs. An AI travel agent, for example, might assemble an itinerary, optimize for loyalty benefits, and monitor for disruptions but still surface choices that require judgment—time versus comfort, cost versus flexibility. A home electronics agent may narrow options based on specifications and reviews but defer to the human when design, compatibility, or brand preference becomes decisive.

In these categories, trust is built not through perfect execution but through explainability and reversibility. As autonomy increases, consumers want to understand not just *what* the agent did but *why* it behaved in that manner. Why did it choose a particular option? Why did it make a substitution? Why did it escalate an exception? Graceful handling of edge cases matters more than success on the happy path.

This is where metadata becomes strategy. Humans infer meaning intuitively, considering factors such as fit, feel, mood, and suitability for a particular occasion. AI agents, of course, do not. They rely on structured, contextual signals. Products that are emotionally legible to people but semantically opaque to machines risk becoming invisible in agent-mediated flows. This requires retailers to invest in rich, machine-readable attributes that enable agents to act with nuance—and to know when to pause and elevate questions to human shoppers.

How value pools shift when agents mediate commerce

Across these patterns, one shift is consistent: the compression of the traditional funnel. Search, comparison, and consideration collapse into a single agent-mediated moment. Continuous commerce replaces episodic decisions. Loyalty becomes less about sentiment and more about policy.

As a result, value pools migrate. Advantage accrues to merchants that can reliably execute against agent constraints, not just those that attract human attention. Margins are shaped by service guarantees, fulfillment reliability, and clarity of policies. For some players, this will unlock efficiency and scale. For others, particularly those dependent on discovery-driven traffic, it introduces the risk of disintermediation.

Importantly, this does not imply a single end state. The automation curve does not prescribe where every category should end up. Instead, it describes the instances where delegation creates value and where it does not. Retailers that recognize these contours early can invest accordingly, pushing toward higher autonomy where it reduces friction and deliberately preserving human moments where they matter most.

The future of commerce is not about maximizing automation. It is about placing autonomy where it enhances experience, economics, and trust. The automation curve offers a practical lens for making those choices. Retailers that use it to guide capability investment, category strategy, and agent readiness will be best positioned to compete as AI agents become an increasingly central interface of commerce.

Deepa Mahajan, **Hannah Mayer**, and **Roger Roberts** are partners in McKinsey's Bay Area office, where **Katharina Giebel** is a consultant. **Katharina Schumacher** is a partner in the Munich office.

The authors wish to thank Eli Stein, Holger Harreis, Jack Trotter, Marcus Keutel, Philipp Kluge, Tara Balakrishnan, and Zamir Lalji for their contributions to this article.

This article was edited by Larry Kanter, a senior editor in the New York office.

This article is part of McKinsey's ongoing research on AI and aims to help business leaders understand the forces transforming ways of working, identify strategic impact areas, and prepare for the next wave of growth. As with all McKinsey research, this work is independent and has not been commissioned or sponsored in any way by any business, government, or other institution. The article and views expressed herein are ours alone.

We welcome your comments on this research at agenticcommerce@mckinsey.com. Learn more about our gen AI insights and sign up for our newsletter.