

Europe's agentic commerce moment: Decision influence is here; execution is coming

A new survey shows that European consumers are using AI to shape *what* they buy, but not *how* they buy—at least not yet.

This article is a collaborative effort by Katharina Schumacher, Marcus Keutel, and Philipp Kluge, with Katharina Giebel and Tiffany Wendler, representing views from QuantumBlack, AI by McKinsey.



Before a European consumer purchases a product, an AI tool may already have narrowed the field to three options—and explained why one of them failed to make the cut.

This is the near-term reality of agentic commerce: AI-mediated discovery and evaluation are scaling quickly, even as full autonomy remains limited. While fully autonomous bots are not buying everything on our behalf, AI is fast becoming the primary interface for discovery, comparison, and recommendation—the point in the consumer journey where preferences form and winners emerge. The stakes are considerable: [McKinsey research](#) estimates that by 2030, agentic commerce could orchestrate \$3 trillion to \$5 trillion globally,¹ as AI agents increasingly influence discovery, decision-making, and transactions across categories. Given Europe's share of global consumption, even moderate adoption scenarios would translate into material impact across European retail and consumer sectors.

For leaders, the strategic question is shifting. It's no longer simply “How do we convert customers?” It's “How do we remain visible and persuasive when the first ‘customer’ in the funnel is not a human, but an AI agent?” When evaluation moves into AI-mediated interfaces, traditional levers such as placement, advertising, and user experience matter less than whether an offering is retrievable, comparable, and defensible in an AI's logic.

New consumer research from France, Germany, and the United Kingdom suggests that Europe is already entering that phase. Most respondents report using AI tools in their everyday personal or professional lives, and usage is not confined to experimentation. Consumers are applying AI to practical decision support: researching topics, synthesizing information, and increasingly shaping purchase decisions. These tools span general-purpose assistants as well as retailer-embedded bots and third-party shopping agents, each with different incentives and economic models.


This article examines how European consumers are already integrating AI into their shopping journeys. Our research shows that AI—whether tools such as ChatGPT and Gemini or brand-owned agents and embedded agentic experiences—is rapidly becoming the primary interface for discovery, comparison, and, increasingly, transaction initiation. It is here, at this early stage of the journey, that preferences are shaped, consideration sets are formed, and competitive advantage is won or lost. Fully autonomous shopping agents, meanwhile, will evolve at different speeds and plateau at different levels across categories and use cases, not always moving all the way to complete, end-to-end autonomy. But the consumer behaviors we are seeing today—such as AI-led comparison, synthesis of options, contextual recommendations, and basket assembly—are leading indicators of deeper delegation to come.

¹ Katharina Schumacher, Roger Roberts, and Katharina Giebel, [The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants](#), McKinsey, October 17, 2025.

Europe's agentic commerce moment: By the numbers

AI is already mainstream in Europe's customer decision journey


Share of European consumers who report using AI tools in their everyday lives, % of respondents



84%


of European consumers report using AI tools in their everyday lives

Use of AI tools by European consumers in their shopping journeys, % of respondents




63%

use AI for comparing options like brands, models, prices, and reviews



55%


use AI to learn about a category or product



46%


use AI to discover or get inspiration for purchases

AI trust in the shopping process for European consumers, % of respondents




56%

are comfortable with AI suggesting options—with the human making final decisions



41%

trust gen AI tools to summarize reviews and highlight trade-offs



47%

are uncomfortable letting AI manage recurring purchases automatically

McKinsey
& Company

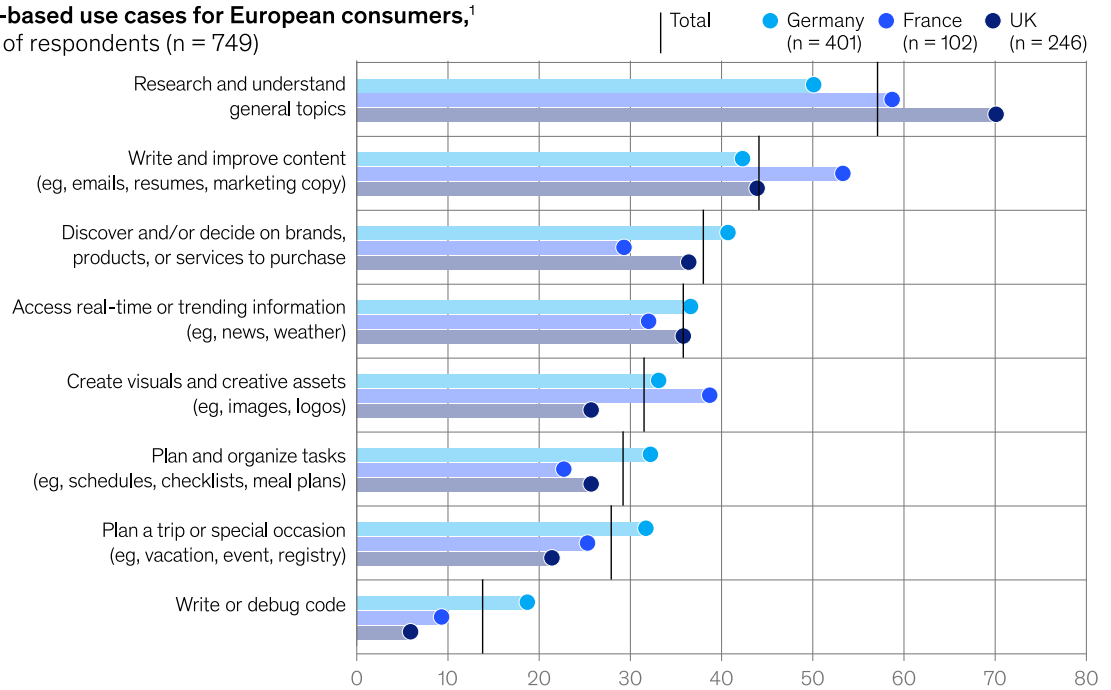
AI is already mainstream in Europe’s consumer decision journey

In a recent McKinsey consumer survey across France, Germany, and the United Kingdom, 84 percent of respondents report using AI tools in everyday life.² The most common use cases include researching general topics (57 percent) and writing and improving content (44 percent). When it comes to researching products and services or deciding what to purchase, 38 percent report using AI tools (Exhibit 1).

Exhibit 1

AI is now mainstream in daily life, with more than a third of consumers in Europe using it to inform shopping decisions.

AI-based use cases for European consumers,¹
% of respondents (n = 749)



¹Question: For which of the following use cases have you used AI tools (eg, ChatGPT, Google AI Overview, Gemini, Perplexity, Microsoft Copilot, Claude, Grok) in the past 3 months?
Source: McKinsey Consumer AI Discovery Survey, Dec 2025

McKinsey & Company

²McKinsey Consumer AI Discovery Survey, December 2025, based on a sample of 749 respondents across France (n = 102), Germany (n = 401), and the United Kingdom (n = 246).

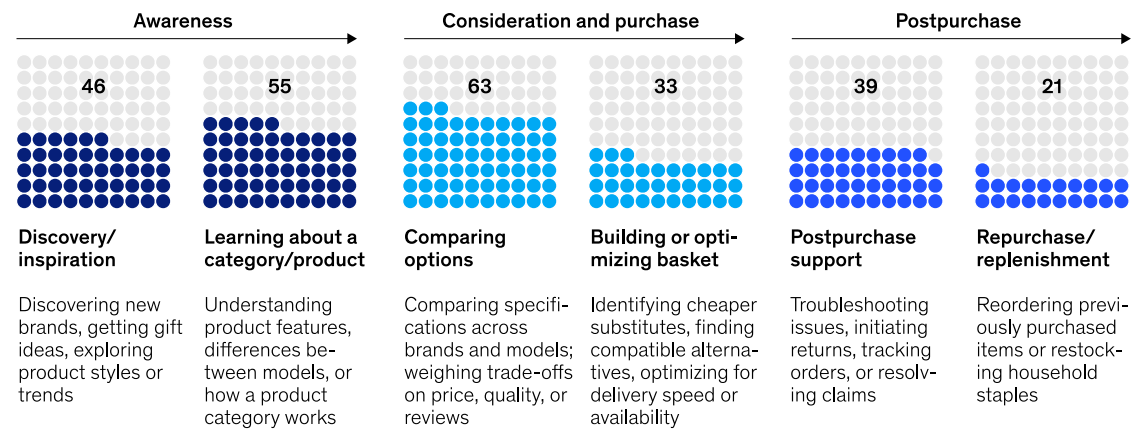
Our data shows that, for European consumers, AI is showing up most strongly upstream of the transaction, at the point where preferences are formed and options are narrowed (Exhibit 2). European consumers most frequently use AI for the following:

- comparing brands, models, prices, and reviews (63 percent)
- learning about a category or product (55 percent)
- discovering new products and getting inspired (46 percent)

Exhibit 2

European consumers primarily use AI upstream in the shopping journey.

European gen AI usage across the shopping journey,¹ % of respondents (n = 749)



¹Question: For which parts of the shopping journey have you used AI tools in the past 3 months?
Source: McKinsey Consumer AI Discovery Survey, Dec 2025

McKinsey & Company

Significantly, usage declines as activities move closer to execution (for example, building baskets and completing checkout). The results point to a clear sequencing effect: AI is being adopted first as a decision-support layer, compressing research, comparison, and synthesis that previously required multiple sites and sources. But the situation is changing fast: Some AI environments are beginning to connect recommendations directly to embedded purchase options, suggesting that conversational commerce may scale before fully autonomous delegation does.

In practice, this means decision influence is scaling faster than execution infrastructure or consumer trust in autonomy. The implication is subtle but significant. By the time a consumer arrives at a retailer or brand interface, much of the evaluative work may already have been completed within AI-mediated workflows or interfaces.

This does not mean execution can be deprioritized: Payments, identity, and authorization are the enabling rails that will determine who captures demand as delegation expands. As agentic capabilities mature, the question will not simply be whether checkout becomes automated, but where, when, and for which categories consumers choose to delegate further. As we explore in our recent article, [“The automation curve in agentic commerce,”](#) delegation is unlikely to expand uniformly; instead, it will scale selectively, shaped by trust, regret risk, and the role human involvement plays in the value of the experience.³ Emerging infrastructure, including open standards for agent-led commerce and payments—such as the Agentic Commerce Protocol (ACP), Agent2Agent (A2A), Agent Payments Protocol (AP2), and Universal Commerce Protocol (UCP)—is beginning to create the rails that could accelerate that shift over time.

Trust concentrates on judgment, not agency

European consumers draw a clear line on how they trust AI across the shopping journey. Trust is highest when AI supports judgment but drops sharply as it moves toward action. Consumers are most comfortable with AI that summarizes reviews, highlights trade-offs, compares options, and recommends a “best” choice—use cases that help people reason better without surrendering control (Exhibit 3). By contrast, trust declines steadily as AI approaches execution—particularly for prefilling baskets, completing checkout, or automatically reordering items.

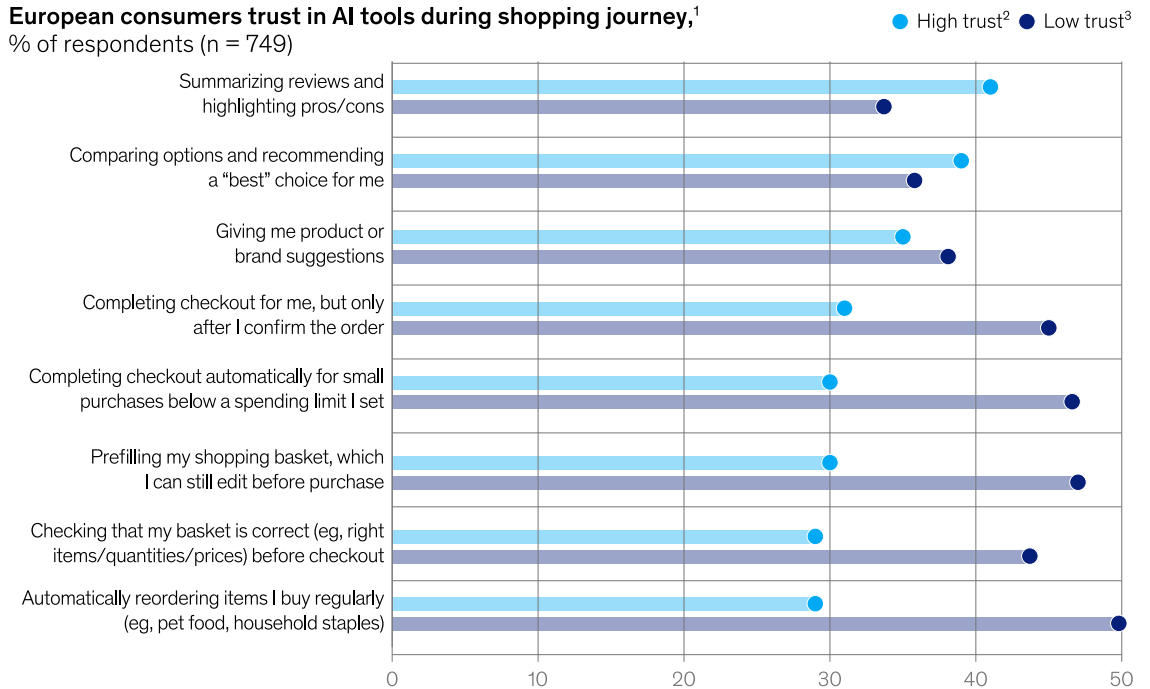
This should not be read as a broad referendum on AI capability. The data suggest something more specific: Consumers are not rejecting AI judgment; they are resisting unbounded authority. Comfort is highest when AI actions are reversible, explicitly authorized, and easy to audit—and lowest when actions are persistent, opaque, or difficult to undo. Seen through this lens, trust is less about whether AI can act and more about whether consumers can see how decisions were made, understand who is accountable, and intervene when needed. In many cases, AI influence is already occurring within conversational interfaces that can embed purchase pathways directly into chat environments, reshaping where discovery and checkout begin. Nonetheless, consumers, by and large, are more willing to outsource cognition than control.

For consumer companies, the strategic question becomes how to design agentic experiences that make reasoning visible, authority explicit, and responsibility clear. Systems that foreground explanation, confirmation, and human override are likely to earn trust faster than those that solely optimize for automation. Adoption will expand first in bounded, episodic use cases—decision support, option narrowing, and assisted checkout—before extending into ongoing or fully autonomous execution. It will be up to organizations to treat trust as a product capability, rather than a communications problem.

³ [“The automation curve in agentic commerce,”](#) McKinsey, January 28, 2026.

Exhibit 3

European respondents trust AI most for evaluative judgment, but far less for autonomous actions.



¹Question: How much would you trust gen AI tools to support you in the following shopping actions?

²"Trust a lot" and "fully trust."

³"Trust a little" and "do not trust at all."

Source: McKinsey Consumer AI Discovery Survey, Dec 2025

McKinsey & Company

Category adoption is broad based, with modest differences by use case

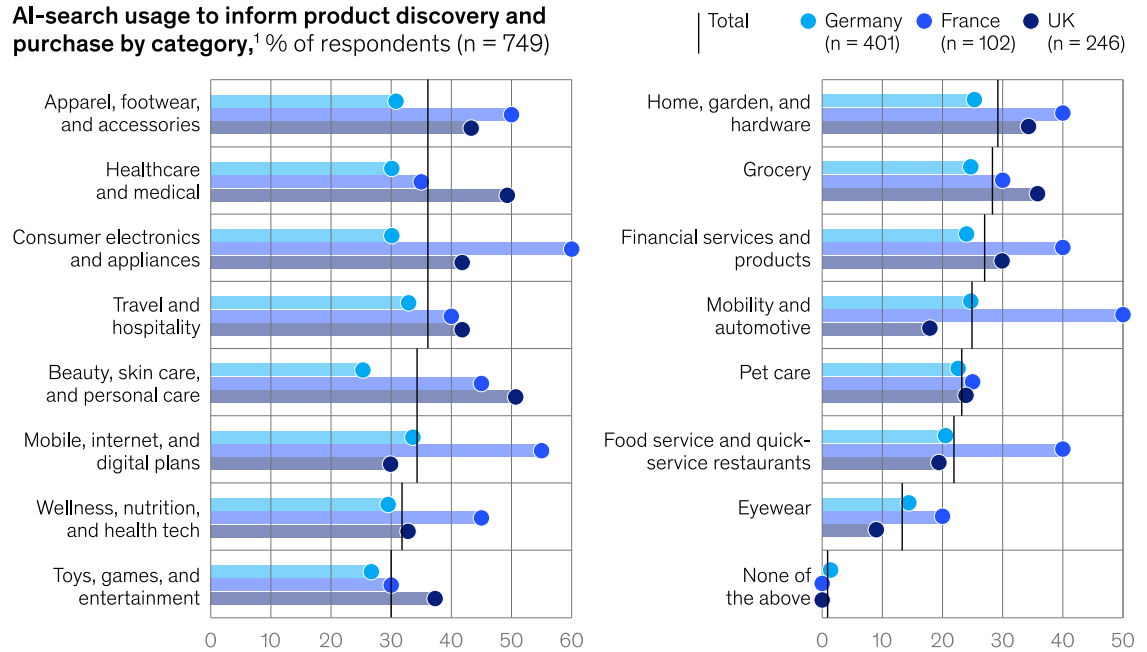
At an aggregate level, the survey reveals a clear, somewhat counterintuitive pattern: Use of AI to inform shopping decisions is broadly distributed across categories, with relatively limited differentiation between most of them. Adoption clusters tightly in the 30 to 36 percent range across a wide set of categories (Exhibit 4), suggesting that AI is already functioning as a general-purpose evaluation layer, rather than a tool confined to a few AI-native verticals.

Usage is slightly higher in research-intensive categories—including apparel, healthcare and medical products, consumer electronics, and travel—where consumers face greater complexity,

Exhibit 4

AI usage in Europe is strongest in research-intensive categories, but broadly adopted overall.

AI-search usage to inform product discovery and purchase by category,¹ % of respondents (n = 749)



¹Question: For which of the following categories have you ever used gen AI tools to discover and decide on brands, products, or services to purchase?
Source: McKinsey Consumer AI Discovery Survey, Dec 2025

McKinsey & Company

trade-offs, or information asymmetry and therefore benefit more from synthesis, comparison, and explanation. However, the differences between leading and midtier categories are modest, reinforcing the idea that AI support is being applied not only in traditionally “complex” purchases but wherever consumers encounter uncertainty or decision friction.

One category stands out clearly: eyewear. At just 13 percent adoption, eyewear lags materially behind the rest of the field. The divergence is notable not because eyewear lacks informational complexity, but because much of the value in eyewear purchasing hinges on physical fit, in-person assessment, and infrequent decision-making—areas where current AI tools offer less incremental value relative to the offline experience.

The eyewear example has important insights for how agentic commerce is likely to scale. Near-term opportunity is shaped less by category labels and more by where AI-driven evaluation creates value without requiring physical validation or sustained delegation. Over time, as interfaces mature and trust boundaries shift, more categories may move up the curve—but adoption is unlikely to progress uniformly or converge quickly on full autonomy.

Questions that will shape how agentic commerce scales in Europe

The European data do not point to a single end state. Instead, they surface a set of unresolved questions that will influence how quickly and how far consumers are willing to delegate decisions to AI systems as agentic commerce matures:

- *How does brand loyalty get expressed when decisions are mediated by agents?* A common concern is that AI-mediated comparison will erode brand loyalty by making switching frictionless. While that outcome is possible in some contexts, particularly for commodity goods and services, the data and emerging agent behaviors point to a more nuanced shift. Brand loyalty is unlikely to disappear, but it may be expressed differently.

As consumers increasingly rely on AI to evaluate and narrow options, brand preferences can enter the decision process through multiple channels. In some cases, consumers may explicitly encode trusted brands or exclusions, effectively providing their agents with “allow lists” that guide recommendations. In others, agents may surface unfamiliar brands that fit a user’s stated preferences, values, or aesthetic, expanding consideration sets while still honoring underlying brand affinities.

Over time, agentic commerce may also introduce new surfaces where brands can assert themselves more directly. As agent-to-agent interactions evolve, brand-specific agents could increasingly act as brand ambassadors—explaining differentiation, negotiating terms, and articulating brand intent within machine-mediated flows. There will also be contexts, such as pure price comparison or routine replenishment, in which brand matters less and functional criteria dominate.

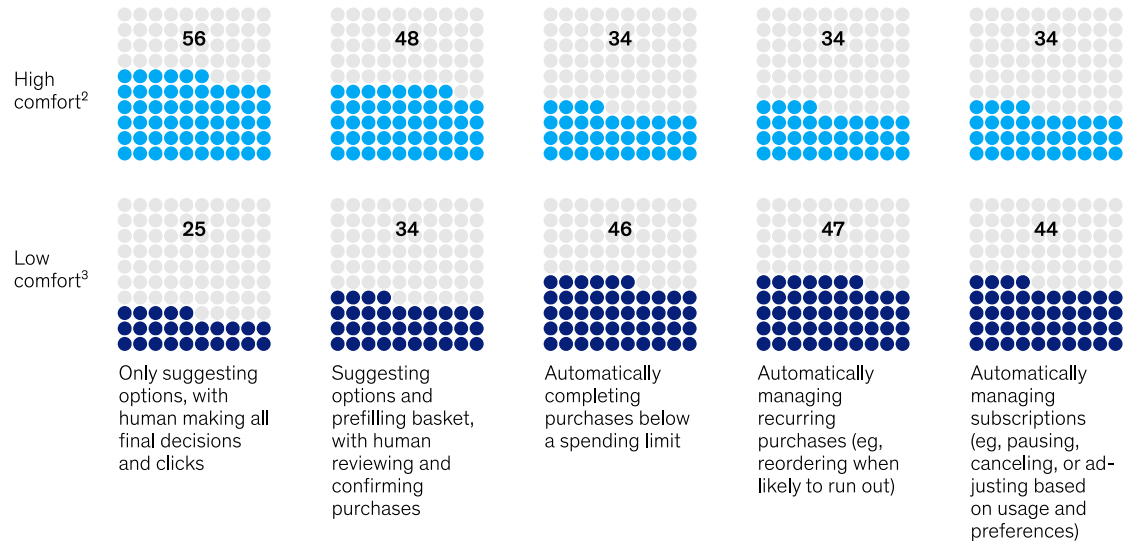
Taken together, the shift suggests not the erosion of brand loyalty, but its reconfiguration from a primarily emotional and visual construct to one that must also be encoded, interpretable, and defensible in agent-mediated decisions.

- *Is Europe’s caution around delegation a constraint or a leading indicator?* Our data show that comfort with AI declines steadily as levels of delegation increase, particularly for persistent or recurring actions such as managing subscriptions or reordering household staples (Exhibit 5). Consumers are most comfortable with AI that assists decision-making while preserving final human control, and least comfortable with AI that operates continuously or without explicit confirmation.

Exhibit 5

Consumers are comfortable delegating judgment—but not ongoing control.

European comfort with AI tool autonomy during the shopping journey,¹ % of respondents (n = 749)



¹Question: For which parts of the shopping journey have you used AI tools in the past 3 months?

²Somewhat comfortable and "very comfortable."

³Somewhat uncomfortable and "very uncomfortable."

Source: McKinsey Consumer AI Discovery Survey, Dec 2025

McKinsey & Company

This pattern could be seen as an expression of conservatism. An alternative, and potentially more consequential, reading is that European consumers are articulating the conditions under which they are willing to delegate, rather than rejecting delegation outright. These conditions include reversibility (the ability to undo actions), accountability (clear responsibility when something goes wrong), and explicit consent (clear boundaries on what the agent is authorized to do).

For leaders, the implication is not to slow investment in autonomy, but to design for conditional delegation—building agentic experiences that earn trust incrementally, make authority explicit, and allow users to calibrate autonomy over time. One test of this hypothesis will be whether agentic solutions that foreground reversibility and explicit mandates see faster adoption, even for more autonomous use cases.

- *Where do consumers draw the line on delegation today?* The comfort sentiment data suggest a clear boundary in how consumers think about delegation. Respondents are most willing to delegate cognitive tasks—such as receiving suggestions, comparing options, or reviewing prefilled baskets—while remaining hesitant to cede ongoing or standing control, particularly for actions that persist over time or operate without repeated confirmation.

This highlights an important distinction between what AI systems are technically capable of and what consumers are currently willing to authorize. Delegation is most acceptable when it is bounded, episodic, and easy to reverse; it becomes more contested as it shifts toward continuous execution or implicit authority. If this pattern holds, adoption of autonomy is unlikely to advance as a single leap toward full automation. Instead, it will expand selectively and conditionally, shaped by how clearly users can define, monitor, and adjust the scope of an agent’s authority over time.

Key imperatives for retailers and brands

The survey findings suggest three important steps that retailers and brands will have to take if they wish to remain relevant in the emerging agentic-commerce era:

- *Compete for AI-mediated decisions (not just human attention).* As discovery and comparison shift into AI interfaces, competitive advantage increasingly depends on whether a brand can be accessed and understood by AI agents, not just discovered and recognized by consumers. Visibility is no longer earned solely through placement, advertising, or user interface; it is earned through machine legibility—what some describe as “agent engine optimization,” in which structured credibility replaces keyword positioning. In practice, this raises the bar on foundational readiness: Rich, structured product metadata, consistent naming and taxonomy, evidence-backed claims, and strong third-party trust signals such as reviews and expert mentions are fast becoming prerequisites for appearing in AI-mediated journeys at all. This is not a wait-and-see moment. It requires strengthening foundational capabilities such as product information management, feed quality, taxonomy consistency, and API exposure, while monitoring and engaging with emerging interoperability standards, such as Model Context Protocol (MCP) and UCP, that aim to structure agent-to-agent commerce. Retailers should begin upgrading their “agent surface area” now—welcoming reputable agent traffic, instrumenting real-time access to accurate data, and differentiating trusted agents from malicious bots. Those that fail to do so risk becoming invisible upstream, before a consumer ever reaches a brand interface.

- *Optimize for explainability, not just conversion.* As consumers increasingly rely on AI to make sense of choices, differentiation is judged less by persuasive user experience and more by clarity, comparability, and proof. In an agentic journey, a value proposition must work in two directions: It must resonate with humans and be structured enough for machines to evaluate, carry forward, and cite in downstream agent-to-agent interactions. Brands that cannot express their differentiation as crisp, evidence-backed “reasons why”—what the product is, who it is for, what trade-offs it makes, and why it fits a given context—risk being filtered out upstream, before a shopper (or their agent) ever reaches a product description page. This creates a new competitive surface for consumer players, one where influence depends not only on shelf presence but also on how effectively attributes and claims can be parsed and compared by AI systems. Winning brands will treat explanation as a product capability: designed, structured, and continuously improved to support both human understanding and machine reasoning.
- *Plan for a web of agents, not a single interface.* Once retailers are visible and explainable in AI-mediated journeys, the next strategic question is how much autonomy to enable, and where. Delegation will expand in stages and is category dependent. Today, as our data show, comfort is highest for bounded, episodic delegation—such as decision support, option narrowing, and assisted checkout—and drops sharply for persistent or standing authority. This has two implications. First, autonomy should be treated as a road map, not a feature switch, starting with reversible, explicitly authorized actions and expanding only as trust, infrastructure, and consumer comfort mature. Second, building a proprietary, consumer-facing AI agent is a selective strategic choice, not a default requirement. It makes sense primarily where a retailer has clear category authority and can add differentiated value through guided exploration, not merely transaction execution.

Businesses should be very clear about the purpose and use cases rather than adopting AI for the sake of adopting AI. Crucially, owning a consumer agent is not the same as owning the customer relationship. In a web-of-agents world—in which personal agents of choice interact with retailer agents or broker agents—retailers can still retain control over the transaction and the relationship by remaining the merchant of record and by owning identity, payment authorization, policies, fulfillment, and post-purchase service. Emerging agentic commerce protocols (ACP, A2A, AP2, and UCP) are explicitly designed to preserve this model, enabling agent-mediated discovery and negotiation while keeping retailers accountable for execution. As a result, competitive advantage shifts away from owning the interface toward being reliably evaluable, negotiable, and executable across agent ecosystems. The winners will be those that design for participation across agent-to-site, agent-to-agent, and brokered models—ensuring their offerings can be understood, trusted, and acted upon by any authorized agent operating on the consumer’s behalf.

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



These consumer signals provide an important demand-side lens on a broader shift already underway. For a deeper perspective on how AI agents are reshaping commerce end to end—including implications for infrastructure, business models, and trust—see McKinsey’s recent report [The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants](#).⁴

Adoption of agentic commerce will be driven by trust

Agentic commerce will likely be defined first by how decision influence reshapes discovery, and only later by how quickly that influence extends into execution. Europe’s consumer data suggest a pragmatic pathway: scale AI-mediated evaluation first, then expand autonomy as enabling rails—identity, authorization, and payments—mature and trust is earned.

That trajectory should not be read as hesitation. It may instead represent a blueprint for trust-aware agentic commerce, where AI earns the right to act by proving it can reason, explain trade-offs, and stay aligned with human intent—before it operates continuously on a consumer’s behalf.

For leaders, the question is no longer whether agentic commerce will matter. It is whether your organization will be visible and differentiated in AI-mediated decisions today—and transaction-ready tomorrow. In a world where agents increasingly speak first, the winners will be those that are interpretable, recommendable, and trusted at every step of delegation.

Katharina Schumacher and **Philipp Kluge** are partners in McKinsey’s Munich office; **Marcus Keutel** is a partner in the Cologne office; **Katharina Giebel** is a consultant in the Bay Area office; and **Tiffany Wendler** is an associate partner in the Stuttgart office.

The authors wish to thank Bony Abraham, Christina Bukow, Holger Harreis, and Roger Roberts for their contributions to this article.

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

This article contributes to 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](#).

Copyright © 2026 McKinsey & Company. All rights reserved.

⁴Katharina Schumacher, Roger Roberts, and Katharina Giebel, [The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants](#), McKinsey, October 17, 2025.