BRAND INTELLIGENCE ARTICLE SERIES

How organizations convert everyday interactions into enterprise intelligence.

Interactions → Data → Learning → Intelligence → Advantage

ARTICLE CODE	BI-AR-04
CATEGORY	Consumer Behavior | AI Transformation
PILLAR	Competing in the Age of AI
CHAPTER	Ch. 3, 5, 7, 9

What Consumers Will — and Won't — Delegate to AI Agents[1]

Baohong Sun

Summary: Every major agentic commerce forecast assumes a consumer who is ready to hand over the buying journey to an AI agent — yet OpenAI quietly pulled direct checkout from ChatGPT after consumers refused to complete transactions through it, even as AI-referred shoppers converted at 31% higher rates on retail sites. The gap is not technological; it is behavioral. This article introduces a Delegation Matrix built on two dimensions most strategies ignore — the cost of choosing versus the value of choosing — and maps four quadrants to four distinct agent modes matched to consumer behavior: the Optimizer for full automation, the Curator for guided selection, the Steward for default management, and the Concierge for experience enrichment. It then develops the trust architecture that moderates delegation across all quadrants and identifies the strategic implications for firms operating in each. The firms building agentic commerce strategies without a behavioral model of what consumers will and won't surrender are investing billions in automating the wrong things.

Keywords: Delegation Matrix; Cost of Choosing; Value of Choosing; Consumer Delegation; Agentic Commerce; Agent Mode; Agent-Trust Equity; Machine Legibility; Preference Learning; Trust Architecture; Intelligence without Delegation; Product-Service-Experience Continuum; User Lifetime Value; Flywheel of Intelligence; Command Center

1. The Delegation Paradox

Every major forecast of agentic commerce assumes a consumer who is ready to hand over the shopping journey to an AI agent. McKinsey projects $3–5 trillion in AI-orchestrated commerce by 2030.[1] Morgan Stanley estimates $190–385 billion in U.S. e-commerce alone.[2] Bain finds that 30–45% of U.S. consumers already use generative AI for product research and comparison.[3] But these projections share a notable gap: none of them specify a behavioral model for *which* decisions consumers will actually delegate — and which they will refuse to surrender regardless of how

capable the technology becomes.

The early evidence suggests this gap matters enormously. OpenAI discontinued direct checkout inside ChatGPT in early March 2026 after low transaction completion rates — consumers were not ready to buy through a conversational agent.[4] Yet Adobe reports that AI-referred shoppers convert at 31% higher rates on retail sites, delivering 254% higher revenue per visit.[5] Both facts are true simultaneously. The first says consumers resist agents *completing* transactions for them. The second says agents are highly effective at *guiding* consumers toward transactions. The difference is not about technology capability. It is about which part of the decision the consumer is willing to hand over — and which part they insist on keeping.

Why should this matter to brands, platforms, and startups building for the agentic economy? Because the answer determines where agent-facing investment generates returns and where it is wasted. A brand that automates stages consumers want to keep will face backlash. A brand that fails to automate stages consumers want to offload will lose to competitors whose agents do. And the boundary between these two zones is not fixed — it varies by category, by consumer segment, by decision stage, and over time as agents learn. Understanding this boundary is the single most important input for any firm's agentic commerce strategy.

BI-AR-02 established what brands must build for agentic commerce: machine legibility, agent-trust equity, and a Command Center that integrates data into real-time intelligence. BI-AR-03 mapped how those requirements differ across the U.S. horizontal and Chinese vertical architectures. This article addresses the prior question that both assume but neither answers: what will consumers actually do? Which decisions will they hand over, which will they keep — and why?

2. The Conventional Wisdom: "Automation Follows Capability"

The prevailing industry view treats consumer delegation as a technology-supply problem. McKinsey's automation curve for agentic commerce exemplifies this framing: it ranks categories by "delegability" based on complexity, risk, and transaction frequency, positioning routine replenishment at the base and high-consideration purchases at the top — implying a natural progression as AI capability climbs.[1] [6] Gartner projects that 40% of enterprise applications will incorporate task-specific AI agents by end of 2026, reinforcing the narrative that delegation expands as capability expands.[7]

This framing falls short in three respects.

First, it ignores that the binding constraint is not agent capability but consumer willingness. Consumers already have the technology to automate grocery replenishment. Amazon's Subscribe & Save, smart refrigerators, and one-click reorder buttons have existed for years — yet only about 23% of U.S. Amazon shoppers maintain even one active subscription.[6] Most consumers do not automate the majority of their grocery purchases — not because the technology fails, but because choosing what to eat, what to cook, what to try this week is part of how people experience daily

life. The technology is ready. The consumer is not — and the reason is not technological.

Second, the capability framing treats delegation decisions as fixed consumer attributes. In practice, they vary by context, mood, time pressure, social setting, and identity salience. The same consumer who delegates commodity restocking on a Tuesday morning may refuse to delegate wine selection for a dinner party on Friday evening. Delegation is not a stable trait. It is a contextual decision made anew each time — and the context includes whether the act of choosing itself carries value.

Third, the conventional view conflates two fundamentally different consumer motivations. Consumers may want an agent to *reduce the cost* of a decision (eliminate information overload, compare alternatives, find the best price). But they may simultaneously want to *preserve the value* of making the decision themselves (express identity, enjoy the process, maintain a sense of control). These are separate dimensions, and confusing them leads to strategies that automate the wrong things.

3. The Consumer Psychology of Delegation

Before introducing a framework, it helps to understand the behavioral forces at work when a consumer decides whether to delegate a purchase decision to an AI agent.

The behavioral economics literature has long established that choosing is not merely instrumental — it is experiential. Iyengar and Lepper's classic work on choice overload showed that too many options can paralyze decision-making, but Schwartz's subsequent research on the paradox of choice revealed something subtler: even when choice is burdensome, people are reluctant to give it up, because the act of choosing reinforces a sense of autonomy and competence.[8] [9] Self-determination theory (Deci and Ryan) provides the deeper explanation: autonomy is a fundamental psychological need, and consumers experience threats to autonomy — including well-intentioned automation — as aversive, even when the automation would improve outcomes.[10] Belk's foundational work on the extended self explains why certain purchases resist delegation entirely: when possessions function as expressions of identity, the act of choosing them is inseparable from their meaning.[11] And Holbrook and Hirschman's research on hedonic consumption established that for many purchase categories, the process of searching, evaluating, and selecting is itself a source of pleasure — not a cost to be minimized but a benefit to be preserved.[12]

These forces do not operate uniformly. In functional categories — insurance, utilities, office supplies — the choosing process carries little intrinsic reward, and consumers experience it as friction. In identity-expressive categories — fashion, gifts, culinary exploration — the choosing process is itself part of the consumption experience. Asking an agent to choose a birthday gift for your partner is not just outsourcing a task; it is surrendering a moment of care that gives the gift its meaning. The key insight is that delegation triggers multiple psychological tensions simultaneously — autonomy, identity, hedonic experience — and the relative strength of each varies by category and context.

The delegation decision is further complicated by the nature of the agent itself. As established in BI-AR-02, agentic commerce introduces two levels of AI: brand-level AI (the Command Center, optimizing for the firm's objectives) and user-level AI (the agentic shopper, optimizing for the consumer's objectives). The consumer's delegation decision is really a question about *which level of AI to trust with which part of the journey* — and whether the agent's objective function genuinely aligns with the consumer's own. Bain's research confirms this asymmetry: consumers trust retail-owned agents three times more than third-party agents, suggesting that institutional credibility transfers to the agent — but also that consumers are acutely sensitive to whose interests the agent serves.[3]

4. The Delegation Matrix: Cost of Choosing vs. Value of Choosing

A more predictive model requires separating two dimensions that the behavioral literature identifies but that industry analysis consistently conflates.

Dimension 1: Cost of choosing. The cognitive, temporal, and informational effort required to make a good decision. It includes information asymmetry (how hard is it to evaluate quality?), alternative proliferation (how many options must be compared?), temporal pressure (how urgent is the decision?), and evaluation complexity (how many attributes must be weighed simultaneously?). When the cost of choosing is high, consumers have a strong incentive to delegate.

Dimension 2: Value of choosing. The intrinsic utility consumers derive from the act of choosing itself — separate from the utility of the product chosen. It includes identity expression, hedonic experience, social signaling, emotional investment, and sense of control. When the value of choosing is high, consumers resist delegation regardless of how capable the agent is.

The intersection produces four distinct delegation patterns — each with a matched agent mode that defines how the AI should behave.

Figure 1: The Delegation Matrix — Cost of Choosing × Value of Choosing

Four quadrants, four agent modes, four strategic responses

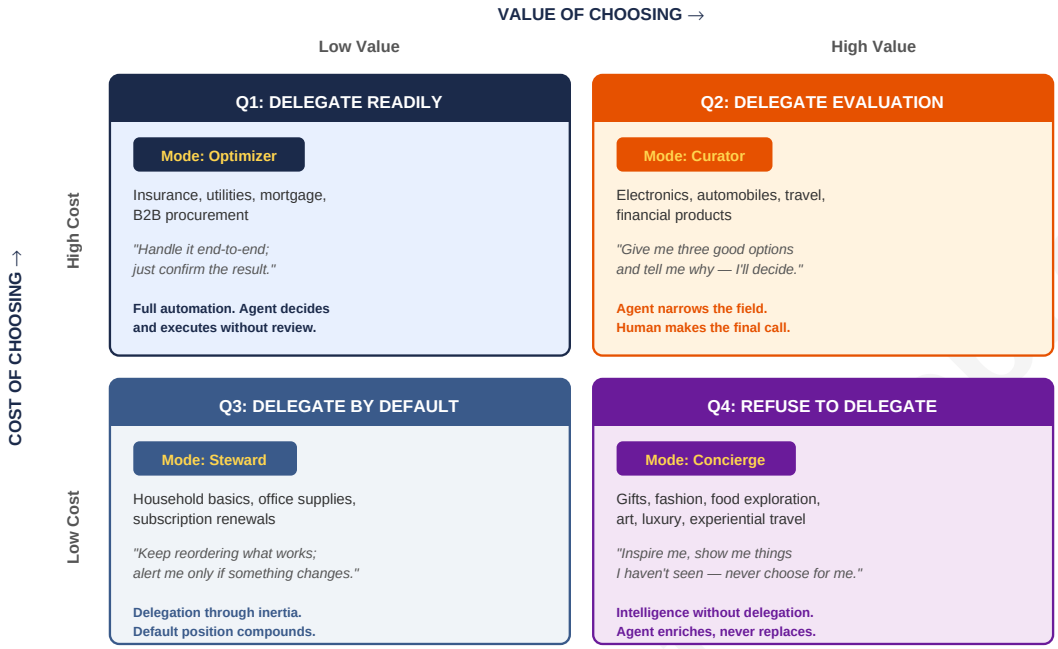

Source: Brand Intelligence (Sun, 2026), Ch. 3, 9. Extended for agentic commerce consumer delegation analysis.

	Low Value of Choosing	High Value of Choosing
High Cost of Choosing	**Q1: Delegate Readily** — Agent mode: *Optimizer*. Insurance, utilities, mortgage, B2B procurement. Consumers want full automation — decisions are cognitively expensive and carry no intrinsic enjoyment. *Consumer expectation:* "Handle it end-to-end; just confirm the result."	**Q2: Delegate Evaluation, Retain Selection** — Agent mode: *Curator*. Electronics, automobiles, travel, financial products. Consumers want agents to narrow the field but insist on making the final choice. *Consumer expectation:* "Give me three good options and tell me why — I'll decide."
Low Cost of Choosing	**Q3: Delegate by Default** — Agent mode: *Steward*. Household basics, office supplies, subscription renewals. Low stakes, low engagement — delegation through inertia. *Consumer expectation:* "Keep reordering what works; alert me only if something changes."	**Q4: Refuse to Delegate** — Agent mode: *Concierge*. Gifts, fashion, food exploration, art, luxury, experiential travel. The act of choosing IS the product. *Consumer expectation:* "Inspire me, show me things I haven't seen — but never choose for me."

Analyzing the Four Quadrants

The four quadrants differ in readiness, market size, and strategic complexity — and these differences determine where agentic commerce investment should be sequenced.

Q1 is the most ready and the easiest to automate. The technology exists, the consumer willingness exists, and the business case is straightforward: reduce friction in decisions that consumers experience as burdens. Insurance comparison, utility switching, and B2B procurement are already seeing agent-mediated transactions at scale (Walmart's Pactum-powered AI negotiation is a Q1 example). The competitive question here is not whether to automate but how to earn the

agent's default selection — which requires machine legibility and protocol presence across the Commerce Control Stack.

Q3 is the easiest to enter but the hardest to defend. Commodity replenishment involves low-complexity decisions that agents can handle immediately. But precisely because the decisions are simple, switching costs are low and differentiation is minimal. The competitive moat comes from post-purchase performance: once an agent learns to trust a brand through consistent fulfillment, the default position compounds through algorithmic switching costs. Programmable deal logic governs who holds that default.

Q2 is the largest market and the most strategically interesting. Electronics, automobiles, travel, and financial products represent enormous transaction volumes — and the hybrid human-agent decision model creates the most complex competitive dynamics. Brands must simultaneously win the agent's shortlist (machine-facing equity) and the human's final selection (consumer-facing equity), while building the outcome record that sustains agent-trust equity over time. Decision-slot competition is the defining battleground.

Q4 is the smallest in agent-mediated volume but the most valuable in brand terms. Luxury, fashion, gifts, and experiences command the highest margins and generate the deepest customer relationships. These categories will resist delegation permanently — not as a transitional state but as a structural feature. The strategic imperative is intelligence *without* delegation: using AI to enrich the choosing experience rather than replace it.

Four Agent Modes

The term "agent mode" rather than "agent type" is deliberate. A single consumer shifts between modes depending on context — behaving as an Optimizer for office supplies, engaging the Curator for electronics, relying on the Steward for household restocking, and seeking the Concierge for a weekend gift — all within the same week. The Delegation Matrix is not only a market segmentation tool; it is a user-state diagnostic. The agent must recognize which mode the consumer is operating in for *this* decision and adjust its behavior accordingly.

The Optimizer is the agent as efficiency engine. It minimizes the cost of choosing entirely — scanning alternatives, evaluating specifications, negotiating terms, and executing transactions without requiring human involvement at any stage. The consumer's expectation is complete automation. The brand's competitive response is to ensure its value is parseable at every layer of the Commerce Control Stack: structured data in the product graph, protocol presence at the transaction rails, and verifiable outcome records that feed agent-trust equity. If the Optimizer cannot read your differentiation, your brand does not enter the consideration set.

The Curator is the agent as intelligent filter. It compresses the evaluation process — synthesizing information, comparing alternatives, eliminating poor options — while preserving the human's authority over the final selection. The Curator must excel at two things simultaneously: machine-facing legibility (to earn the decision slot from the agent ecosystem) and consumer-facing

resonance (to present differentiation in terms the human can *feel* at the moment of selection). The brand that wins Q2 is the one whose advantages are both machine-parseable and emotionally compelling — verifiable by the agent and desirable to the human. Winning the Curator's shortlist is necessary but not sufficient; the brand must also win the human's heart at the moment of choice.

The Steward is the agent as default manager. It maintains continuity — reordering what works, managing subscriptions, alerting the consumer only when something changes. The competitive challenge is holding the default position against algorithmic switching: once the Steward learns to trust a brand through reliable fulfillment, that trust creates algorithmic switching costs. The brands that lose Q3 are those with inconsistent fulfillment — because the Steward's memory is perfect and its tolerance for unreliability is zero.

The Concierge is the agent as experience enhancer. Its job is not to choose for the consumer but to make the act of choosing richer — surfacing unexpected discoveries, personalizing the browsing experience, providing context and inspiration, and enabling customization. This is where the product-service-experience continuum from Chapter 7 of the Brand Intelligence framework becomes the strategic guide. In Q4, the brand's opportunity is to expand along the dimensions of service and experience — using intelligence to transform a product purchase into a multi-layered experience that deepens engagement. A luxury watch brand does not need an agent that selects watches. It needs an agent that learns the consumer's aesthetic sensibility, curates personalized discovery journeys, facilitates virtual try-on, and connects the buyer to the brand's community of collectors. The intelligence works *behind* the choice, enriching the consumer's autonomy rather than replacing it.

5. The Trust Dimension

The Delegation Matrix maps delegation based on the economics of choosing — cost versus value. But human decision-making is not purely economic. In a world where the agent that advises you may also serve the brand that sells to you, behavioral factors like trust, perceived alignment, and transparency become even more decisive than in human-to-human commerce. An imperfect human salesperson can build rapport, read body language, and earn trust through interpersonal signals. A machine agent cannot — which means trust must be designed into the system architecturally, not earned through charisma. This makes the trust dimension not a secondary consideration but a structural prerequisite for delegation in every quadrant.

Consumer trust in AI agents operates on three levels.

Competence trust — "Can the agent do this well?" This is the dimension most industry analysis addresses. It rises predictably as agents demonstrate capability. Bain's research confirms that familiarity accelerates trust: 72% of consumers have used AI tools in some form, but only 10% have completed a purchase through AI — suggesting the competence gap is closing faster than the willingness gap.[3]

Alignment trust — "Is the agent working for me or for the brand that deployed it?" This is the more consequential barrier. When a brand's agent recommends the brand's own product, the consumer cannot easily distinguish between genuine suitability and self-serving optimization. The Brand Intelligence framework identifies this as the core tension of human-machine symbiosis (Ch. 9) — the degree to which algorithmic optimization aligns with human welfare determines whether the system earns or erodes trust. Agent-trust equity is built precisely by demonstrating alignment over time through verifiable outcomes.

Transparency trust — "Can I understand what the agent did and why?" Delegation requires that the consumer can audit the agent's reasoning — at minimum, after the fact. The absence of interpretability does not prevent delegation in low-stakes contexts (Q3), but it effectively blocks delegation in high-stakes contexts (Q1 and Q2) where the consequences of a poor decision are significant. Objective-function governance — understanding what agents optimize for and making brand value visible to that logic — is the brand's response to the transparency challenge.

Figure 2: Trust Requirements by Quadrant and Agent Mode

Three trust levels moderate delegation willingness differently in each quadrant

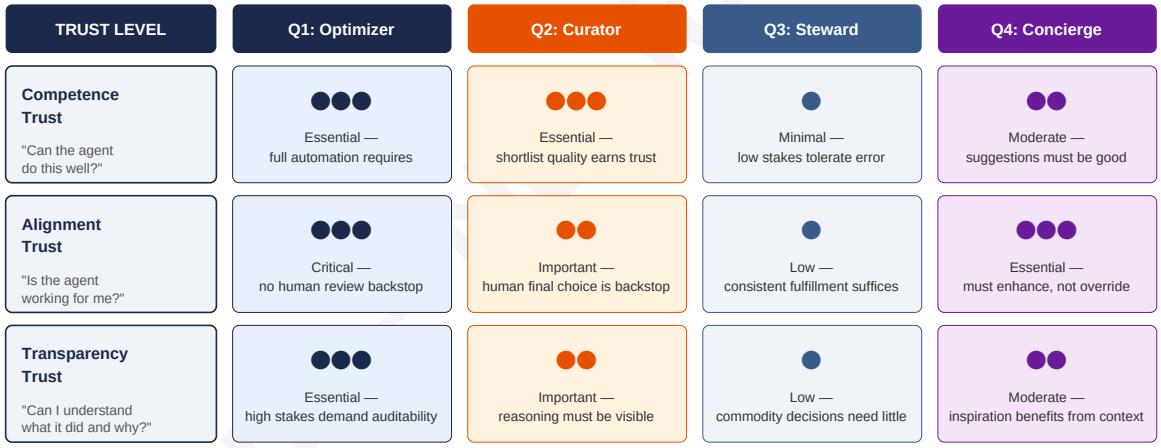

Source: Brand Intelligence (Sun, 2026), Ch. 3, 9. Trust architecture for consumer delegation in agentic commerce.

Trust requirements differ by quadrant and by agent mode. Q1 demands all three levels: consumers are handing over consequential decisions entirely, so the Optimizer must demonstrate competence, alignment, and transparency simultaneously. Q2 requires strong competence and alignment trust at the evaluation stage, but the human retains the final decision as a trust backstop — the Curator earns trust by the quality of its shortlist and the clarity of its reasoning. Q3 needs only minimal competence trust — the stakes are low enough that the Steward's consistent fulfillment record is sufficient. Q4 has the most interesting trust dynamic: here, trust is not about delegation at all but about whether the Concierge's suggestions *enhance* the consumer's own judgment without overriding it. Understanding these trust requirements leads directly to the question of strategic response — how firms in each quadrant should build for the agentic economy.

6. Strategic Implications

The strategic challenge differs fundamentally by industry. A health insurer competing in Q1 faces an entirely different agentic commerce problem than a luxury fashion house in Q4 — yet most industry strategies treat agentic commerce as a single adoption curve that all firms ride at different speeds. The Delegation Matrix reframes the question: the strategic response must be organized around the quadrant the firm actually occupies, the agent mode its consumers operate in, and the trust architecture that enables delegation in that context.

Quadrant	Primary Strategic Challenge	Key Capabilities to Build	Immediate Priority
Q1: Delegate Readily (insurance, utilities, B2B)	Earning the agent's default selection — the brand the agent chooses without human review	Machine legibility across every Commerce Control Stack layer; protocol presence (UCP, MCP); programmable deal logic for automated negotiation	Structure every offer as agent-parseable data; invest in GEO; treat operational consistency as demand generation
Q2: Delegate Evaluation, Retain Selection (electronics, auto, travel, financial)	Earning the agent's shortlist while winning the human's final choice	All three forms of brand equity simultaneously: consumer-facing, machine-facing, and agent-trust equity	Design decision-ready offers; invest in post-purchase outcome capture; make non-price differentiation machine-verifiable
Q3: Delegate by Default (household, office, subscriptions)	Holding the default position against algorithmic switching	Post-purchase performance that compounds through the Flywheel of Intelligence; algorithmic switching costs	Maximize fulfillment consistency; build first-party data loops via the Put-and-Take Method; design auto-replenishment that captures learning
Q4: Refuse to Delegate (luxury, fashion, gifts, art, experiences)	Deploying intelligence without replacing the choosing experience	Consumer-facing equity; community as anti-delegation moat (Ch. 5); product-service-experience continuum (Ch. 7)	Invest in experiential infrastructure; cultivate the Concierge mode; expand along service and experience dimensions

A critical design question for firms: should they build one agent or multiple agent modes? The answer is that the most effective systems will deploy different modes — or different behaviors within the same agent — depending on the consumer's context. The consumer's position in the decision journey (awareness, research, evaluation, purchase, post-purchase) shifts the cost-value calculus within each quadrant. A Q2 consumer at the research stage welcomes Curator initiative; the same consumer at the purchase stage insists on control. Firms that design their agents to adjust behavior across both the quadrant dimension (what kind of decision) and the journey dimension (what stage of that decision) will outperform those that deploy a static agent personality.

The Agent as Digital Twin

There is a deeper architectural implication that connects the Delegation Matrix to the Brand Intelligence framework. As the agent interacts repeatedly with a consumer — learning preferences,

constraints, and aspirations across multiple modes — it becomes something more than a shopping tool. It becomes a digital twin of the consumer: a computational representation that mirrors the consumer's decision-making patterns with increasing fidelity over time. The Optimizer learns what "lowest cost" means for *this* consumer. The Curator learns which tradeoffs *this* consumer will accept. The Steward learns *this* consumer's tolerance for substitution. The Concierge learns *this* consumer's aesthetic sensibility and emotional triggers.

The same logic that the Brand Intelligence framework applies to the human user applies to this digital twin: users generate data, data trains algorithms, algorithms improve experiences, better experiences attract more engagement. The Flywheel of Intelligence (Ch. 2, 10) does not distinguish between the human user and the machine representation; it compounds from both. The ULTV framework extends naturally: the digital twin generates data value (behavioral patterns), expresses preferences (a form of social value), and transacts (monetary value). The brands that recognize this — that treat the consumer's agent as a user in its own right, worthy of its own relationship strategy — will build intelligence advantages that compound across both human and machine interactions.

7. Forward Look

The Delegation Matrix describes the behavioral terrain as it exists today — but that terrain is not static. Three developments will reshape the boundaries between what consumers delegate and what they keep.

Delegation boundaries will move — but not uniformly. As agents prove reliable, some consumers will shift from Curator mode to Optimizer mode in specific categories. A consumer who today insists on choosing their own running shoes may, after three successful agent-selected pairs, stop caring. But the movement is not one-directional. New categories of high-value-of-choosing will also emerge — AI-curated experiences, personalized products, co-created goods — that resist delegation by design. The Delegation Matrix is not static; it evolves with both technology and culture.

Agents will develop learned preference models that blur delegation and expression. In machine learning and recommendation systems, the process by which an agent learns a consumer's preferences is called *preference learning* — the computational construction of a model that represents the consumer's aesthetic sensibility, ethical commitments, quality standards, and identity-expressive patterns. The resulting representation — a *preference embedding* — encodes a consumer's taste in a form that enables the agent to act on preferences the consumer has never explicitly stated.[13] When a personal agent selects a wine that matches the consumer's learned taste profile, is that delegation or expression? If the agent has internalized the consumer's aesthetic identity so deeply that its selections feel like the consumer's own, the boundary between surrendering a decision and extending one's judgment dissolves. This hybrid mode — the agent as an extension of identity rather than a replacement for judgment — may become the dominant mode in Q2 categories, fundamentally reshaping what decision-slot competition looks like. The firms that master preference learning at the individual consumer level will hold the most defensible

position in agentic commerce, because a well-trained preference model is a switching cost that no competitor can replicate without the same depth of interaction data.

Heterogeneous and dynamic consumer behavior is the strategic reality. The firms that will thrive are those that design for the full spectrum — Optimizer, Curator, Steward, and Concierge modes — rather than assuming a single adoption trajectory. Incumbent brands must protect the post-purchase relationships that feed agent-trust equity while making pre-purchase stages machine-legible. Platforms must design for delegation heterogeneity — enabling both full automation and human-in-the-loop experiences within the same interface. Startups have an opportunity to serve the underserved quadrants: building intelligence-enhanced choosing tools for Q4 categories, or creating alignment-transparent agents that solve the trust gap in Q2.

The deeper lesson of this analysis is behavioral, not technological. The firms building agentic commerce strategies based on technology roadmaps — "as agents get smarter, consumers will delegate more" — are building on sand. The firms that start with the consumer — understanding the psychology of autonomy, identity, trust, and the irreducible human desire to choose — will build on bedrock. The Delegation Matrix is not a prediction of where agentic commerce will end up. It is a map of the behavioral terrain that will determine how it gets there. The terrain is human. The firms that forget this will automate the wrong things. The firms that remember it will build the agentic economy that consumers actually want.

References

1. McKinsey, "The Agentic Commerce Opportunity," October 2025. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-agentic-commerce-opportunity-how-ai-agents-are-ushering-in-a-new-era-for-consumers-and-merchants
2. Morgan Stanley, "Agentic Commerce Impact Could Reach $385 Billion by 2030," Morgan Stanley Research, December 2025. https://www.morganstanley.com/insights/articles/agentic-commerce-market-impact-outlook
3. Bain & Company, "Agentic AI Commerce Hinges on Consumer Trust," 2025. https://www.bain.com/insights/agentic-ai-commerce-hinges-on-consumer-trust/
4. Digital Commerce 360, "OpenAI shifts checkout plans in its agentic commerce strategy," March 6, 2026. Modern Retail, "Shopify says purchases are coming 'inside ChatGPT' through agentic storefronts as OpenAI retreats on Instant Checkout," March 2026.
5. Adobe, "Holiday Shopping Season Drove a Record $257.8 Billion Online," press release, January 2026. https://news.adobe.com/news/2026/01/adobe-holiday-shopping-season
6. McKinsey, "The Automation Curve in Agentic Commerce," 2025. Subscription adoption data: Consumer Intelligence Research Partners, Amazon Prime Member Survey, 2025.
7. Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," press release, August 26, 2025.
8. Iyengar, Sheena S., and Mark R. Lepper. "When Choice Is Demotivating: Can One Desire Too Much of a Good Thing?" *Journal of Personality and Social Psychology* 79, no. 6 (2000): 995–1006.
9. Schwartz, Barry. *The Paradox of Choice: Why More Is Less.* New York: Ecco, 2004.

10. Deci, Edward L., and Richard M. Ryan. "Self-Determination Theory: A Macrotheory of Human Motivation, Development, and Health." *Canadian Psychology* 49, no. 3 (2008): 182–185.
11. Belk, Russell W. "Possessions and the Extended Self." *Journal of Consumer Research* 15, no. 2 (1988): 139–168.
12. Holbrook, Morris B., and Elizabeth C. Hirschman. "The Experiential Aspects of Consumption: Consumer Fantasies, Feelings, and Fun." *Journal of Consumer Research* 9, no. 2 (1982): 132–140.
13. On preference learning and preference embeddings: Fürnkranz, Johannes, and Eyke Hüllermeier, eds. *Preference Learning.* Berlin: Springer, 2011. For application in recommender systems: Koren, Yehuda, Robert Bell, and Chris Volinsky. "Matrix Factorization Techniques for Recommender Systems." *Computer* 42, no. 8 (2009): 30–37.
14. Sun, Baohong. *Brand Intelligence: Navigating the Transformation in the AI and Web3 Era.* Springer Nature, 2026. https://link.springer.com/book/10.1007/978-3-032-17490-6

Cross-references: *Brand Intelligence* Ch. 2, 3, 5, 7, 9, 10 | BI-AR-02: The Agent Shopper — Brand Strategy for the Agentic Commerce Era | BI-AR-03: The Agent Divide — Why Agentic Commerce Looks Different in the U.S. and China | BI-CS-03: NIO (community and advocacy)