BRAND INTELLIGENCE ARTICLE SERIES

How organizations convert everyday interactions into enterprise intelligence.

Interactions → Data → Learning → Intelligence → Advantage

ARTICLE CODE	BI-AR-02
CATEGORY	AI Transformation | Brand Strategy
PILLAR	Competing in the Age of AI
CHAPTER	Ch. 2, 3, 4, 7, 8, 9, 10

The Agent Shopper: Brand Strategy for the Agentic Commerce Era[1]

Baohong Sun

Summary: Most executives think of agentic commerce as a smarter chatbot. It is not. A new class of AI shopping agents — built overwhelmingly by third parties, not by brands — is rewriting how purchase decisions are initiated, evaluated, and executed. These agents carry their own objective functions, decision logic, and learning systems; they develop what this article calls agent intelligence from inputs that brands do not control. This article introduces machine legibility — the degree to which a brand's value is structured, accessible, and interpretable by AI agents — as the organizing concept for competing in this new landscape. It maps who is building the agents and what constitutes their commercial intelligence, analyzes how the agent's five core decisions reshape the traditional User Journey, and explains why the brand's existing intelligence architecture already serves both human and machine users through different interfaces. The strategic shifts that follow are not what most executives expect — and the brands unprepared for them face not visible decline but invisible exclusion from the agent's consideration set. The firms that have invested in their Brand Intelligence architecture are better positioned than they realize. The firms that have not are running out of time.

Keywords: Agentic Commerce; Agent Shopper; Machine Legibility; Agent Intelligence; Decision-Slot Competition; Objective-Function Governance; Composability Competition; Deal-Space Competition; Confidence-Model Competition; Agent-Trust Equity; Verifiable Value; Programmable Deal Logic; Flywheel of Intelligence; Put-and-Take Method

1. The Agent Shopper Has Arrived

The pace of change is striking. In December 2025, Klarna launched its Agentic Product Protocol — an open standard that made over 100 million products across twelve markets instantly discoverable by AI agents.[1] Weeks later, Google announced the Universal Commerce Protocol (UCP — an open standard for agent-merchant interoperability) at the National Retail Federation

conference, backed by more than twenty partners including Shopify, Walmart, Target, Mastercard, and Visa.[2] OpenAI and Stripe co-developed the Agentic Commerce Protocol (ACP), an open standard for AI-enabled payments and commerce — though OpenAI's initial attempt at direct checkout inside ChatGPT was quietly discontinued in early March 2026 after low transaction completion rates, a revealing signal that consumer purchasing behavior in agent environments is still forming.[3] Perplexity expanded its AI shopping assistant to more than 5,000 merchants through PayPal integration.[4] And in a revealing legal confrontation, a federal court in California granted Amazon a preliminary injunction against Perplexity in March 2026, finding that AI agents accessing Amazon's systems on behalf of users did not constitute authorized access by the retailer's standards.[5]

These are not incremental upgrades. They represent the emergence of a new commercial architecture in which a growing share of buying decisions are evaluated, filtered, and sometimes executed by systems operating on behalf of human buyers — what this article calls the **agent shopper**. The scale projections are substantial: Gartner projects that 40% of enterprise applications will incorporate task-specific AI agents by the end of 2026, up from less than 5% in 2025.[6] Adobe reported a 693% increase in traffic to U.S. retail sites from generative AI tools during the 2025 holiday season, with AI-referred shoppers converting at 31% higher rates and delivering 254% higher revenue per visit.[7] Morgan Stanley estimates that agentic commerce could reach $190–385 billion in U.S. e-commerce spending by 2030.[8] McKinsey frames the global opportunity at $3–5 trillion in AI-orchestrated commerce revenue by the same year.[9] For brand leaders, these figures raise a question that decades of marketing theory did not prepare them for: what happens to brand strategy when the buyer is not a person but an agent shopper acting on a person's behalf?

2. Who Builds the Agent Shoppers — and What Constitutes Their Intelligence

Before asking what brands should do, executives need to understand who is building the agent shoppers — because most of them are not being built by brands. Third parties define the agent's intelligence; brands must respond.

The Builders

Five categories of firms are constructing the agent ecosystem:

AI platforms and assistants (OpenAI, Google, Apple) control conversational interfaces and sit at the moment of intent capture. Google's UCP connects Gemini and AI Mode in Search directly to merchant backends.[2] These platforms are building the general-purpose agent that interprets what the user wants and orchestrates the path to fulfillment.

Commerce infrastructure providers (Stripe, payment networks, identity services) build the transaction plumbing beneath the agent's decisions — how purchases are authorized, routed, and completed.[3]

Retailers, marketplaces, and merchant ecosystems are building their own agent-compatible systems. Klarna's protocol makes merchant catalogs machine-readable.[1] Walmart deploys its own Sparky assistant while simultaneously integrating with Google's Gemini and OpenAI's ChatGPT — a deliberate strategy to secure decision slots across multiple agent ecosystems.[10]

Aggregators and data providers organize supply, comparison, and review data at scale, serving as normalized, trusted sources of commerce data for agent shoppers. Bazaarvoice aggregates product reviews and user-generated content across thousands of brands — the structured outcome data that agents use for trust calibration. Syndigo and Salsify provide product data syndication platforms that normalize product attributes across retailers, directly enabling the machine legibility that agent shoppers require for evaluation and comparison. Feedonomics optimizes product feeds for algorithmic consumption. These firms occupy a critical intermediary position: they shape what the agent can see and how it compares alternatives, even though the underlying data originates with brands.

Brands themselves — but selectively. The pattern of digital history suggests caution: consumers have consistently migrated to general-purpose tools that serve *their* interests. But some brands will build narrower agents around high-consideration categories, concierge service, replenishment, or post-purchase guidance.

What Constitutes Agent Intelligence

An AI shopping agent is not a blank search engine. It carries a set of core components that together constitute what we call **agent intelligence** — the full set of design parameters, functional capabilities, and learning inputs that determine how an agent shopper decides. Understanding these components is critical because the question of who influences each determines who holds power in agentic commerce.

The Brand Intelligence framework (Sun, 2026) provides the conceptual parallel. The Command Center — the brand's algorithmic decision engine — generates real-time marketing decisions by optimizing *who* to engage, *what* to offer, *when* to act, *where* to reach them, and *how* to present it, all to maximize User Lifetime Value (ULTV) over the entire customer relationship. The consumer's agent shopper solves a symmetric problem: who to buy from, what to buy, when to buy, where to transact, and how to execute — optimizing for whatever objective the user defines. The seven components below are what give the agent shopper the capability to solve this problem. As shown in Figure 1, they fall into three layers — Architecture (the agent's design parameters), Capabilities (its functional reach), and Learning (the inputs that make it smarter over time):

Agent Architecture — the design parameters that define *who the agent is*. These are relatively stable, set by the platform and user before the shopping process begins. The AI platform (Google, OpenAI, Apple) designs the reasoning logic and sets default heuristics; the user expresses preferences and sets constraints — though many users will accept platform defaults they never explicitly chose.

Objective function. What the agent optimizes for: lowest price, lowest regret, maximum long-term satisfaction, sustainability, or some weighted combination. An agent optimizing for "best value" shops very differently from one optimizing for "fastest delivery."

Decision logic. How the agent weighs evidence, handles uncertainty, and resolves trade-offs among alternatives. When Google's agent and OpenAI's agent evaluate the same category, they may reach different conclusions not because the data differs but because the reasoning architecture differs.

Constraints and guardrails. Hard boundaries the agent must respect: budget ceilings, brand exclusions, dietary or health restrictions, ethical red lines, regulatory requirements. Unlike the objective function (what to optimize *for*) or preferences (soft signals), constraints are non-negotiable filters that remove options before evaluation begins.

Agent Capabilities — the functional reach that determines *what the agent can see and do*. These are shaped by protocols, data providers, and the brand's own digital infrastructure within the brand-owned ecosystem. The brand has the most direct influence here: it determines the completeness and quality of its product data, the depth of its digital commerce infrastructure, and the protocols through which agents can interact with its systems. Aggregators and platforms curate what the agent actually sees — but the underlying data originates with the brand.

Product graph and structured data. What the agent can see and compare: the completeness, accuracy, freshness, and machine-readability of product information, pricing, inventory, service terms, and reviews across agent-facing feeds.

Transaction capabilities. What the agent can *do*: browse a catalog, check real-time inventory, configure a product, place an order, process a return. The protocols (UCP, ACP, MCP) and the brand's own systems define these capabilities.

Agent Learning — the inputs that feed the agent and make it *smarter over time*. These evolve with every interaction. Learning is a shared domain: the user feeds preference signals into the agent's learning loop, while the brand feeds the post-purchase outcomes — fulfillment quality, service experiences, product performance — that shape the agent's future behavior. Over time, the brand that consistently delivers accumulates a compounding advantage in the agent's learning layer.

Preference and context inputs. The user's purchase history, stated preferences, household constraints, brand loyalties, and past satisfaction signals — sometimes provided explicitly, sometimes inferred from behavioral patterns observed over time.

Learning from outcomes. The agent accumulates experience from past transactions — fulfillment accuracy, return rates, customer satisfaction signals — and adjusts its future behavior accordingly. Brands that consistently deliver feed the agent a strengthening trust signal. Brands that disappoint are quietly routed around.

Together, these seven components constitute **agent intelligence** — the commercial mind of the agent shopper. The three-layer classification matters strategically: Architecture determines which brands can compete for the agent's attention. Capabilities determine which brands the agent can actually evaluate and transact with. Learning determines which brands the agent develops loyalty to over time.

Figure 1: The Agent Intelligence Stack

Seven components across three layers that determine how an AI shopping agent decides

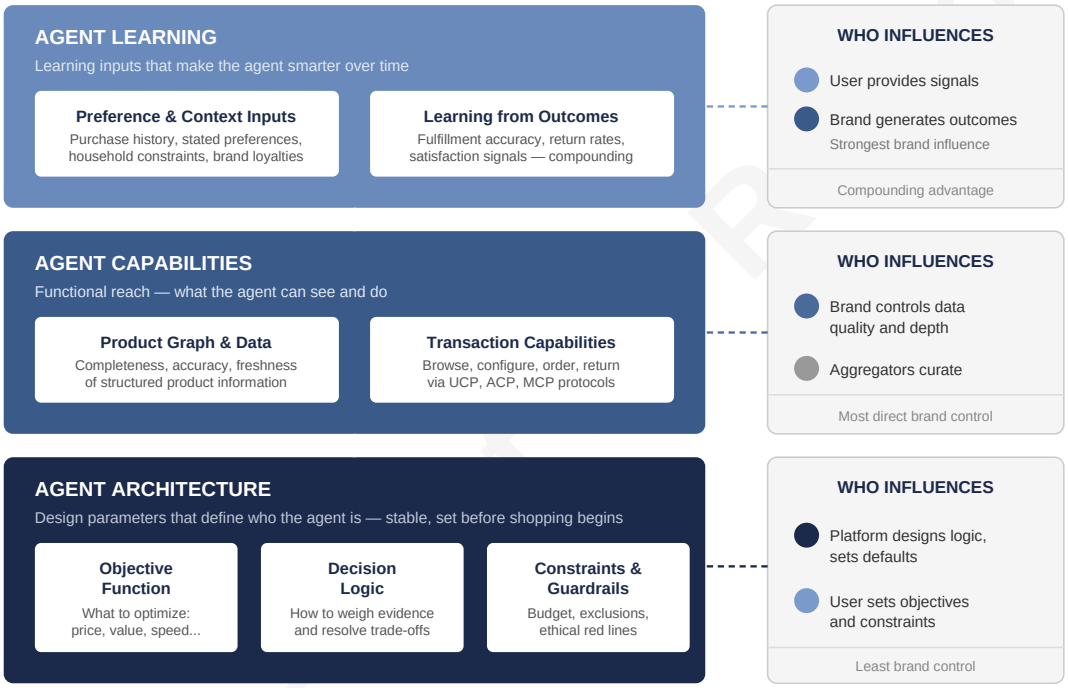

Source: Brand Intelligence (Sun, 2026), Ch. 9. Adapted for agentic commerce.

Machine Legibility: The New Strategic Imperative

This analysis — who builds the agent shoppers, what constitutes their intelligence, and who influences each layer — points to a single organizing concept for brand strategy in the agentic era: **machine legibility**.

Machine legibility is the degree to which a brand's value — its products, pricing, service quality, differentiation, and trustworthiness — is structured, accessible, and interpretable by AI agents. It is not simply having a website or a product feed. It is making the brand's entire value proposition readable and actionable by systems that discover, evaluate, configure, transact, and learn. Machine legibility must also work across modalities: a brand whose value exists only in visual storytelling will be invisible in voice-based agent interactions; a brand whose differentiation requires a long narrative will lose in environments where the agent presents three-line summaries. The requirement is format-agnostic legibility — conversational, visual, and embedded simultaneously.

3. The Five Decisions an Agent Shopper Makes — and What Brands Must Feed Them

What does machine legibility actually mean in practice? The answer maps to five types of decisions — and for each, the brand must ask whether it provides what the agent needs to decide in its favor.

Discovery: "Does this brand belong in my consideration set?" This is a binary gate — inclusion or exclusion — based on category relevance, attribute match, and minimum data completeness. If the brand doesn't pass this gate, nothing else matters. The legibility requirement is *findability*: standardized taxonomy, structured attributes, presence in agent-facing feeds, and Generative Engine Optimization (GEO) readiness — the nine tactics for making a brand's knowledge base optimized for AI retrieval and generation (Sun, 2026, Ch. 10).

This is where *decision-slot competition* occurs — the struggle to be included in the agent's shortlist before the consumer ever sees the options. In traditional digital commerce, brands competed on the search page or in the feed. In agentic commerce, the competition happens inside the agent's reasoning process, and the losers are never seen by the consumer.

Evaluation: "How does this rank against alternatives?" The agent scores and compares across specifications, price, reviews, service terms, and reputation signals. This is where the objective function bites — the same alternatives produce different rankings depending on whether the agent optimizes for cost, quality, sustainability, or minimal regret. The legibility requirement is *comparability*: standardized attributes in formats agents can score, transparent pricing that includes total cost of ownership, and machine-readable differentiation.

Here *objective-function governance* becomes a strategic issue. The brand cannot control which platform the consumer uses or how that platform's default logic interprets an ambiguous instruction like "find me the best value." But it can ensure that its advantages are legible to *whichever* interpretation prevails — price, satisfaction, risk minimization, or sustainability. The brands whose differentiation is only *asserted* in advertising copy will lose to brands whose differentiation is *evidenced* in structured, verifiable data.

Configuration: "What specific option best fits this user?" The agent selects the variant, bundle, add-ons, service level, and delivery option that match the user's constraints and preferences. This goes beyond ranking to constructing the actual offer. The legibility requirement is *modularity*: product data structured so the agent can explore combinations, compatibility rules, personalization parameters, and bundle logic.

This is where *composability competition* emerges. The brand that offers rigid, fixed SKUs loses to the brand whose offer architecture is modular enough for the agent to construct the optimal combination for a specific user. The competition is not about having the best single product — it is about exposing the richest solution space. A brand with ten configurable parameters gives the agent more degrees of freedom to find a fit than a competitor with a take-it-or-leave-it listing. In

agentic commerce, the most composable offer wins the configuration decision.

Transaction: "At what terms, when, and how?" The agent decides timing, channel, payment method, and — in more advanced interactions — negotiated terms (price, warranty, bundle trade-offs, delivery commitments). The legibility requirement begins as *actionability* (real-time inventory, delivery windows, payment protocol compatibility) and evolves toward *negotiability*: programmable deal parameters, price floors and ceilings, concession logic, escalation triggers.

This is the arena of *deal-space competition*. Brands no longer compete solely on the price they set — they compete on the parameter space within which their systems can negotiate. One brand publishes a fixed price and waits. Another exposes programmable deal logic: price floors and ceilings, bundle trade-offs, loyalty incentives, delivery flexibility, warranty upgrades. The agent will gravitate toward the richer negotiation space because it can optimize across more dimensions simultaneously. The wider and more intelligent the deal space, the more likely the agent finds terms that beat a competitor's static offer.

Trust calibration: "How much do I trust this brand, and how do I update that trust?" This cuts across every other decision. The agent assigns a confidence weight to each brand based on accumulated signals and updates that weight after every interaction. The legibility requirement is *verifiability*: observable data the agent can incorporate into its confidence model. This is what builds *agent-trust equity* over time — and it compounds.

Unlike the other four decisions, trust calibration creates *confidence-model competition* — a form of rivalry that accumulates rather than resets. Every fulfilled promise strengthens the brand's position in the agent's confidence model; every broken promise weakens it. Unlike human trust, which is fuzzy and forgetful, agent trust is cumulative and precise. The brand with a longer, cleaner performance record earns a structural advantage that new entrants or inconsistent performers cannot quickly overcome. This makes operational excellence not just an efficiency metric but a compounding competitive asset.

The quality test for machine legibility is this: *for each of the five decisions, can an agent shopper — with no prior knowledge of your brand — find what it needs to decide in your favor?* Most brands today are partially legible for Discovery (they have some structured data) and nearly illegible for everything else.

How Agent Decisions Reshape the User Journey

The five agent decisions map to, but do not replicate, the five stages of the traditional User Journey: need recognition, information search, evaluation, purchase, and post-purchase experience. Figure 2 illustrates three structural transformations that change how brands must compete.

Figure 2: How Agent Decisions Reshape the User Journey

Five UJ stages map to five agent decisions through three structural transformations

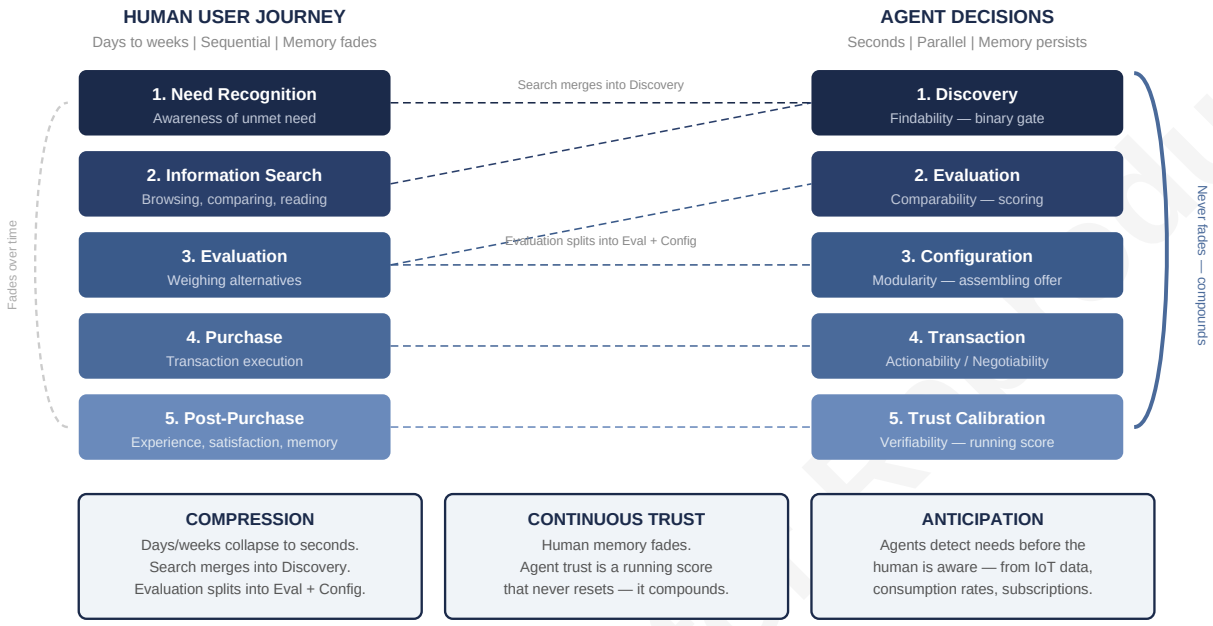

Source: Brand Intelligence (Sun, 2026), Ch. 3. User Journey stages adapted for agentic commerce.

Stages compress. What takes a human days or weeks, the agent executes in seconds. Information search collapses into Discovery. The human evaluation stage splits into Evaluation and Configuration — because agents explore variant and bundle combinations programmatically in ways humans rarely do. The strategic window for brand influence shrinks from days to milliseconds: machine legibility must be immediate.

Trust becomes continuous. In the human User Journey, post-purchase feeds back loosely through memory, which fades. The agent's trust calibration is a running score that never fades — every fulfillment success, every return handled smoothly informs *all* future decisions. This makes trust calibration the only agent decision that is truly cross-cutting, shaping every other decision continuously rather than operating as a terminal stage.

Need recognition becomes anticipatory. The agent can monitor usage patterns — consumption rates from smart product data, subscription cycles, seasonal patterns — and initiate purchases before the human is aware of the need. Brands must be legible not just when the customer is searching but *before the search begins*.

These three transformations — compression, continuous trust, and anticipation — redefine what a brand must build. The question shifts from *what message to send* to *what infrastructure makes the brand legible to an agent shopper that decides in milliseconds*.

The Trust Signal Stack

Consumer delegation to agents raises its own trust dynamics — competence, alignment, and transparency — which BI-AR-04 examines through the Delegation Matrix. Here we focus on the

other direction: how the agent builds trust in the brand.

The fifth agent decision — trust calibration — determines how agents update their confidence in a brand over time. Trust is not a single metric. It draws on multiple categories of verifiable signals, and not all of them require a prior transaction:

Aggregated outcome data — review scores, return rates, complaint volumes, satisfaction metrics compiled across all users and agents. The machine equivalent of peer trust: crowd-sourced, comparative, and available before the agent's first transaction with the brand.

Third-party verification — machine-readable certifications, safety ratings, independent test results (such as those from consumer research and quality ranking firms), regulatory compliance records. Authority-based trust signals that exist independently of any transaction history.

Data quality and transparency signals — how complete, consistent, and honest the brand's structured data is. A brand that provides full total-cost-of-ownership data, transparent return policies, and accurate inventory signals behaves in a trustworthy way in the pre-purchase interaction itself. The agent can infer trustworthiness from how the brand presents its information.

System reliability — does the brand's API respond consistently? Are protocol interactions stable and fast? This technical trust signal accumulates with every query, not just with completed transactions.

Direct post-purchase outcome data — fulfillment accuracy, delivery speed, dispute rates, return friction, service quality, and subscription retention from the agent's own transactions with the brand. This is the most powerful trust signal because it is first-party, proprietary, and compounding — the agent's direct experience, not someone else's report.

Of these five signal types, direct post-purchase outcome data is strategically the most important. It is the one competitors cannot replicate simply by structuring their catalog better. A connected product that generates continuous usage data — driving patterns, charging behavior, software interaction — transforms a one-time transaction into a longitudinal intelligence asset. The brand that feeds this depth of outcome data into the agent's learning loop earns trust that a brand whose relationship ends at the point of sale cannot match.

The strategic implication: operational inconsistency is no longer just an internal efficiency problem. In agentic commerce, it becomes a demand problem — because the machine remembers every failure, and its memory does not fade.

4. The System That Serves Both — If You Built It

The five agent decisions — discovery, evaluation, configuration, transaction, trust calibration — each require the brand to make a specific dimension of its value machine-legible: findability for discovery, comparability for evaluation, modularity for configuration, actionability and negotiability for transaction, and verifiability for trust calibration. Delivering all five requires

infrastructure. The good news: brands that have invested in the Brand Intelligence architecture already possess much of what agent shoppers need. They need to open a second door.

One Architecture

A single intelligence ecosystem — accumulating users, data, and algorithms as strategic assets — is organized around eight foundational modules, three technical layers, and four information flows. Five types of brand touchpoints feed this ecosystem: online communities generating peer-driven content and advocacy data, brand-owned online stores serving as the intelligence and revenue core, smart products producing continuous usage signals through IoT-enabled interfaces, smart stores capturing physical-world behavioral data through digitally enhanced retail environments, and mobile apps functioning as the super touchpoint and super transmitter that orchestrates the User Journey across all channels. All five feed data to the Command Center, which integrates it into real-time marketing intelligence through the Flywheel of Intelligence: users generate data, data trains algorithms, algorithms improve experiences, and better experiences attract more users (Sun, 2026, Ch. 4–9).

This architecture was designed to transform marketing from campaign execution into a continuously learning system. It now turns out to serve a second purpose: the same infrastructure that generates intelligence for human users simultaneously makes the brand readable for machine users. The investment was not wasted; it was preparation.

Two Users

The intelligence ecosystem now serves two types of users: human shoppers and their agent shoppers. Both generate data that feeds the Flywheel. Both are served by the same Command Center logic. But they interact through fundamentally different interfaces.

The human interacts through the **UX Center** — visual, emotional, experiential interfaces delivered through apps, stores, communities, and smart products. The human browses, feels, compares, hesitates, and decides through a mix of reason and emotion. The brand's investment in user experience, storytelling, and sensory design serves this user.

The agent shopper interacts through a **Protocol Layer** — REST APIs (Representational State Transfer — the standard web protocol for data queries), MCP endpoints (Model Context Protocol — the standard for connecting AI agents to external tools and data), A2A negotiation channels (Agent-to-Agent — the protocol for autonomous agent-to-agent communication) — that exposes the same underlying intelligence in structured, machine-readable form. The agent queries, scores, configures, transacts, and learns through data and logic. The brand's investment in structured data, commerce infrastructure, and post-purchase intelligence serves this user.

Figure 3 illustrates this dual-interface architecture — the same Brand Intelligence ecosystem at center, with the UX Center serving human users and the Protocol Layer serving machine agents, both feeding the same Flywheel of Intelligence.

Figure 3: One Architecture, Two Users, Three Tiers

The Brand Intelligence ecosystem serves humans and machines through different interfaces

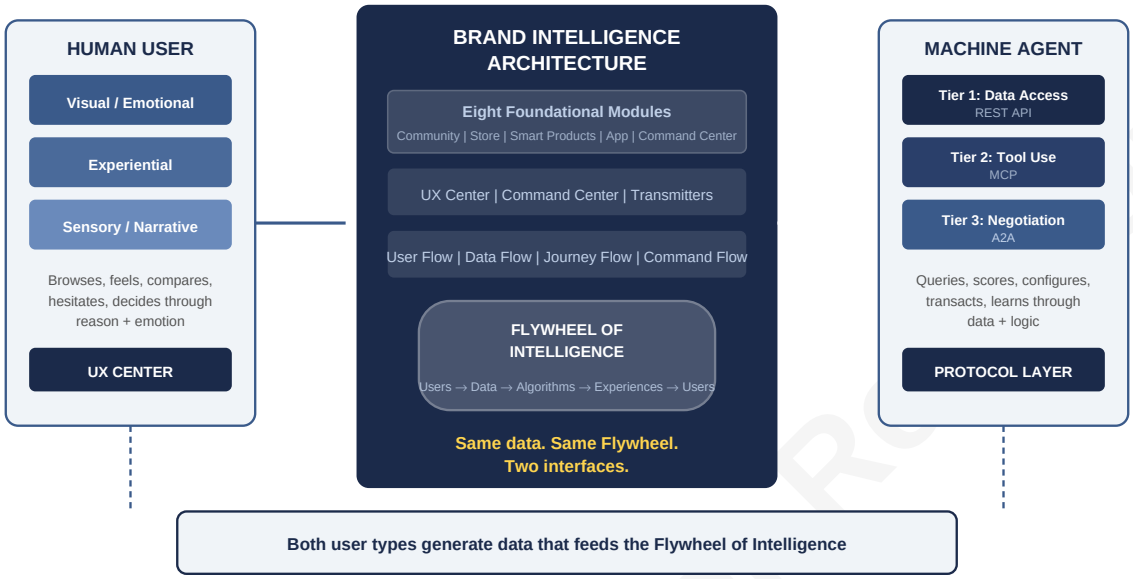

Source: Brand Intelligence (Sun, 2026), Ch. 4. Extended for agentic commerce dual-user architecture.

Same ecosystem. Same data. Same Flywheel. Two interfaces.

Three Tiers of Interaction

Just as brands build UX touchpoints — stores, apps, communities, smart products — to serve human shoppers, they need to build three types of programmatic interface for agent shoppers. The protocol layer operates at three progressively deeper levels — and understanding this architecture is essential for every strategic decision that follows.

Tier 1: Data Access (REST API). The agent queries the brand's system for structured product data — pricing, inventory, specifications, availability, reviews. This is a catalog model: the brand publishes, the agent reads. No real-time interaction, no negotiation. The brand's role is to be *discoverable*. If the agent cannot read your data, you do not exist in its consideration set.

Tier 2: Tool Use (MCP). The agent uses the brand's commerce system as a tool — browsing catalogs, checking real-time inventory, configuring products, placing orders. In its Summer 2025 Edition, Shopify activated a default MCP endpoint on every store on its platform, making merchant catalogs automatically accessible to AI agents without requiring any developer setup. MCP is more interactive than REST: the agent can *act* within the brand's system, not just read from it. The brand's role is to be *transactable*.

Tier 3: Autonomous Negotiation (A2A). Two autonomous agents — one representing the buyer, one representing the brand — communicate directly, exchanging offers, counteroffers, constraints, and reaching agreements at machine speed. The brand's role is to deploy *negotiation intelligence* — the capability to negotiate autonomously while protecting margin and preserving the customer

relationship. This tier is still emerging in consumer commerce but already operational in B2B procurement.

These three tiers are not alternatives — they are layers. A brand needs Tier 1 readiness to be discovered, Tier 2 readiness to be transacted with, and Tier 3 readiness to compete when agents begin negotiating on behalf of consumers. The five agent decisions from Section 3 map directly onto these tiers: Discovery and initial Evaluation operate at Tier 1; Configuration and Transaction operate at Tier 2 and Tier 3; Trust calibration operates across all three. Table 1 summarizes the five decisions, their legibility requirements, and primary tier.

Agent Decision	Legibility Requirement	Competitive Arena	Primary Tier
Discovery — Does this brand belong in my consideration set?	Findability: standardized taxonomy, structured attributes, GEO readiness, protocol presence	Decision-slot competition: invisible exclusion — losers are never seen	Tier 1
Evaluation — How does this rank against alternatives?	Comparability: standardized formats, transparent pricing, machine-readable differentiation	Objective-function governance: whose definition of "best" controls the ranking	Tier 1–2
Configuration — What specific option best fits this user?	Modularity: modular product data, configurable offers, agent-explorable combinations	Composability competition: richest solution space wins the assembly	Tier 2
Transaction — At what terms, when, and how?	Actionability → Negotiability: programmable deal logic, concession rules, escalation triggers	Deal-space competition: widest negotiation bandwidth wins the terms	Tier 2–3
Trust calibration — How much do I trust this brand?	Verifiability: observable outcome data that feeds the agent's confidence model	Confidence-model competition: cumulative performance creates compounding advantage	All tiers

Table 1: The Five Agent Decisions, Legibility Requirements, and Competitive Arenas

This is also where brand ownership becomes strategically critical. The brand cannot own the agent's Architecture — that belongs to the platform. But it can own the data that feeds Agent Capabilities and the outcomes that feed Agent Learning. The brands that invest in structured product data, agent-accessible digital infrastructure, and post-purchase intelligence are building assets that compound across every agent interaction. The brands that outsource these to intermediaries are ceding the components of agent intelligence they could have owned.

5. How Agentic Commerce Changes Brand Strategy

What Brands Are Competing For

The infrastructure described in Section 4 — one architecture serving two users through three tiers of interaction — is the foundation. But what are brands actually competing for in this new landscape? Five competitive arenas — one for each agent decision — define the contest, as mapped in Figure 4.

The decision slot. Before a consumer sees any option, the agent shopper has already filtered, ranked, and shortlisted. The competition for inclusion in that shortlist — decision-slot competition — happens inside the agent's reasoning, invisible to brands that lose. There is no "page two" of results. There is no shelf the consumer walks past. The agent simply never considers you.

The objective function. The same set of alternatives produces different rankings depending on whose definition of "best" governs the comparison. Objective-function governance — the question of who controls this definition when the consumer's instructions are ambiguous — determines which brand the agent favors. The brand cannot control the platform's default logic, but it can ensure its advantages are legible to whichever interpretation prevails.

The solution space. The agent does not just pick a product — it assembles the best version of a product from available components. Composability competition rewards the brand that exposes the richest set of configurable options: variants, bundles, service levels, add-ons. The brand with more degrees of freedom gives the agent more ways to find the optimal fit.

The deal space. As agent-to-agent negotiation matures, the brand that can only publish a price will lose to the brand that can negotiate across multiple parameters — at machine speed, across thousands of simultaneous conversations. Deal-space competition shifts the unit of rivalry from the price point to the negotiation architecture (BI-AR-05 develops this concept in detail).

The confidence model. Trust calibration is a running score that compounds over time. Every successful fulfillment, every accurate delivery, every honest product description strengthens the brand's position in the agent's confidence model. Confidence-model competition is a long game — won through operational consistency, not marketing campaigns — and it creates an incumbent advantage that new entrants cannot quickly replicate.

Each of the five shifts below addresses one or more of these competitive arenas.

Figure 4: Five Strategic Shifts in Agentic Commerce

What changes when agents choose — from the world brands know to the world they must build

Shift	FROM (Human-Centric World)	TO (Agentic World)	Key Concept
1	**Storytelling** Narrative, visual, emotional — designed for humans who feel	**Structured Intelligence** + machine-readable, verifiable — both languages simultaneously	Machine Legibility
2	**One Brand Equity** Consumer-facing only — associations, trust, loyalty in human minds	**Three Equities** Consumer-facing + machine-facing + agent-trust equity (compounding)	Agent-Trust Equity
3	**Asserted Differentiation** Brand says it is better — advertising, claims, positioning	**Evidenced Differentiation** Data shows it is better — verifiable warranties, outcomes, TCO	Verifiable Value
4	**Static Pricing** Set the right price, publish it, wait for buyers	**Programmable Deal Logic** Design the negotiation space — agents negotiate at machine speed	Deal Architecture
5	**Platform Dependence** Reach through intermediaries at the cost of data sovereignty	**Put-and-Take Sovereignty** Interoperate for reach, then convert to brand-owned relationships	Put-and-Take Method

Shift 1: Speak Two Languages

For decades, marketing has been the art of persuasion — storytelling, emotional resonance, sensory experience, aspirational imagery. These capabilities built the world's most valuable brands, and they are not going away. But agentic commerce introduces a parallel requirement: the ability to communicate value in ways that machines can read, compare, and act upon.

The brand must now speak two languages simultaneously. The first is narrative — visual, emotional, experiential — designed for human shoppers who decide through a mix of reason and feeling. The second is structured — rational, procedural, verifiable — designed for agent shoppers that process data, follow decision logic, and learn from measurable outcomes. A brand that speaks only the first language will remain powerful with humans but invisible to the agents that increasingly mediate access to those humans. A brand that speaks only the second will be legible to machines but meaningless to the people whose preferences the machines are supposed to serve. The brands that master both languages will compound advantage on both sides of the Flywheel. The brands that speak only one will lose the other audience.

Shift 2: Three Brand Equities

Brand equity used to mean one thing: what the brand is worth in the minds of human consumers. In agentic commerce, it means three.

Consumer-facing equity is what brand management has always built: the associations, emotions, trust, and loyalty that make a customer willing to pay more, choose faster, and forgive more readily. Consumer-facing equity drives the choices that humans refuse to delegate — the birthday gift, the anniversary trip. Nothing about agentic commerce diminishes its importance in these contexts.

Machine-facing equity is the brand's legibility to intelligent systems — the five legibility requirements working together. A brand with strong consumer-facing equity but weak machine-facing equity faces a specific danger: the human customer may prefer it, but the agent may never present it as an option. This is the invisible loss — the brand does not see a decline in conversion rates because the customer never arrives.

Agent-trust equity is earned machine trust built through repeated positive outcomes that the agent observes and weights in future decisions. It is the cumulative record that shapes how an agent shopper behaves toward the brand over time — built entirely through verifiable signals, not narrative. Agent-trust equity compounds: every successful fulfillment, every accurate delivery, every honest product description strengthens the brand's position in the agent's confidence model.

The strategic danger is asymmetry. A brand can be strong in one form of equity and fatally weak in another — and the weakness may be invisible because the excluded brand never knows it was excluded. The executive's task is to audit all three and invest accordingly.

Shift 3: Prove It — The Verifiable Value Imperative

When agent shoppers can compare alternatives instantly across price, delivery, service conditions, and post-purchase outcome data, weak differentiation becomes easier to expose. The deeper question is not whether agents will compress prices. It is how the concept of value itself changes when the buyer is a machine.

Verifiable value is differentiation that can be measured, compared, and confirmed through data. A warranty that is machine-readable. A fulfillment record that is checkable. A satisfaction rate that is observable. This is the distinction between *asserted* differentiation (what the brand says about itself) and *evidenced* differentiation (what the data shows). In agentic commerce, evidenced differentiation holds pricing power. Asserted differentiation does not.

This is why *objective-function governance* matters. The brand cannot control which platform the consumer uses or how that platform interprets "best value." But a brand whose warranty coverage, service quality, and total cost of ownership are structured in machine-readable formats can defend margin whether the agent optimizes for price, for satisfaction, or for risk minimization. The brand's pricing power in an agentic world is directly proportional to the share of its differentiation that is verifiable rather than merely asserted.

Shift 4: Deals That Negotiate Themselves

Today, most discussions of agentic commerce focus on search, recommendation, and checkout. The next phase will involve negotiation among agents — the Tier 3 interaction where buyer-side and seller-side systems dynamically adjust bundles, terms, service levels, and incentives before a transaction is completed.

If recommendation determines who enters the shortlist, negotiation may determine who wins the transaction. The unit of competition shifts from static offers to *programmable deal logic* — the full set of parameters within which the brand's agent is authorized to negotiate: price floors and ceilings, bundle configurations, service-level trade-offs, loyalty incentives, delivery options, warranty terms, return conditions, and escalation triggers. The strategic capability shifts from pricing strategy (setting the right price) to deal-space design (architecting the negotiation parameter space that maximizes expected value across all counterparty configurations). The brand that can only publish a price will lose to the brand that exposes a richer deal space — at machine speed, across thousands of simultaneous conversations.

Shift 5: The Sovereignty Trap

Agentic commerce amplifies the sovereignty risk that brands face with both human and machine shoppers — losing the customer relationship to an intermediary. To remain visible to agent shoppers, brands must become interoperable with external protocols and platforms. But if those intermediaries become the primary interface, the brand risks losing control over demand formation, customer insight, and the Flywheel of Intelligence — the self-reinforcing loop in which users generate data, data trains algorithms, algorithms improve experiences, and better experiences attract more users.

The Put-and-Take Method (Sun, 2026, Ch. 10) provides the operational response. The method describes a two-phase strategy: brands use third-party platforms and agent ecosystems for broad reach and discovery (the "put"), then engineer the post-interaction experience to convert anonymous, agent-mediated interactions into identified, permissioned relationships within the brand-owned ecosystem (the "take"). The customer registers for a product profile. The product ships with a connected diagnostic tool. The follow-up sequence generates behavioral data no external agent can access.

For agent shoppers, Put-and-Take is both harder and more important than for human shoppers. Harder because the intermediary platform may prevent the brand from capturing the customer's identity — the agent may complete the transaction without ever passing personal information to the brand. More important because without the "take," every sale through an external agent becomes a one-time event rather than a compounding intelligence asset. The brands that solve this problem will build intelligence advantages that compound with every cycle. The brands that do not will become fulfillment layers inside someone else's system.

6. What Executives Should Do Now

Agentic commerce will not reshape every market at the same pace. But across all categories, it is a strategic issue, not a narrow e-commerce experiment. The work ahead maps directly to the three tiers, five agent decisions, and five strategic shifts developed in this article — organized as three layers that build on each other, as shown in Figure 5.

Figure 5: The Three Action Layers — Executive Roadmap

Each layer builds on the one beneath. Start from the Foundation.

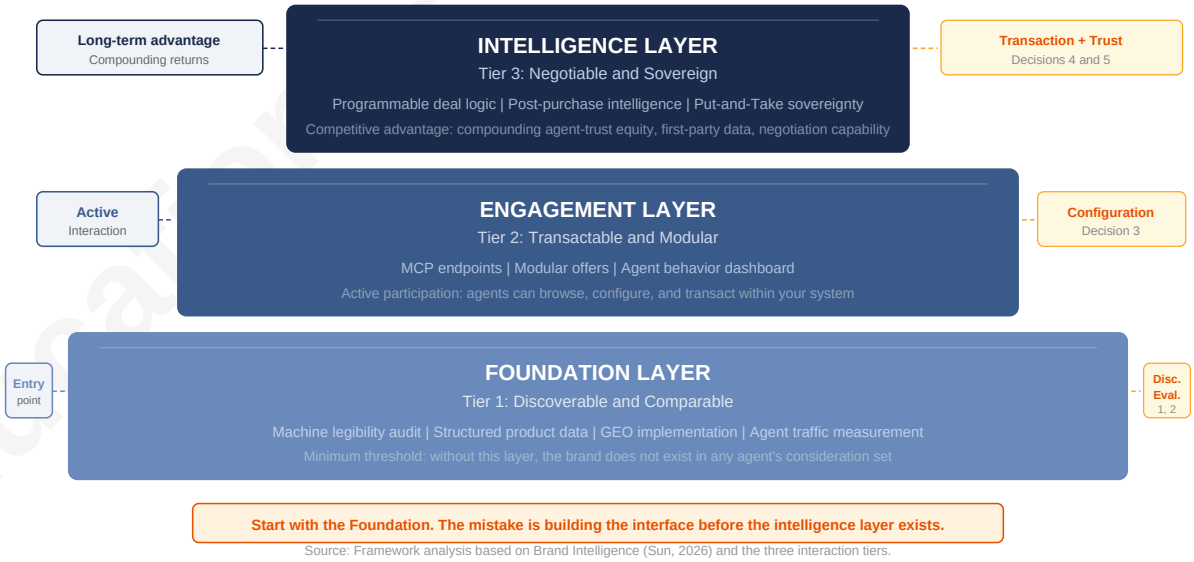

Source: Framework analysis based on Brand Intelligence (Sun, 2026) and the three interaction tiers.

Foundation Layer: Be Discoverable and Comparable (Tier 1)

Every organization starts here. Without Tier 1 readiness, the brand does not exist in any agent's consideration set.

Audit machine legibility across the five decisions. Most brands will find they are partially legible for Discovery and invisible for everything else. The audit should map current data completeness, format standardization, and protocol presence against each of the five legibility requirements.

Structure product data for agent consumption. Make attributes, pricing (including total cost of ownership), service terms, warranties, and return policies machine-readable across all agent-facing feeds. This is the minimum threshold for decision-slot competition.

Implement GEO. The nine Generative Engine Optimization tactics (Sun, 2026, Ch. 10) are the Tier 1 foundation of machine legibility — ensuring the brand's knowledge base is optimized for AI retrieval and generation.

Measure agent-influenced traffic. Begin tracking visits, queries, and transactions originating from AI agents — identifiable through user-agent strings, API access logs, and protocol-specific request patterns. What you cannot measure, you cannot improve.

Engagement Layer: Be Transactable and Modular (Tier 2)

With the foundation in place, the brand becomes an active participant in the agent's decision process — not just readable but interactable.

Activate agent-accessible endpoints. Ensure commerce infrastructure supports MCP or equivalent protocol access — enabling agents to browse, configure, and transact within the brand's system in real time.

Design modular, decision-ready offers. Structure products, bundles, service levels, and delivery options so agents can explore combinations programmatically — not static listings but configurable offers the agent can assemble for specific users.

Build an agent behavior dashboard. Track which agents query the brand's systems, what they ask for, how they evaluate alternatives, and which offers they select or reject. The Command Center's CDP should begin ingesting agent interaction data alongside human behavioral data — because agent behavior is now a first-class data source.

Intelligence Layer: Prepare for Negotiation and Sovereignty (Tier 3)

The most advanced layer builds long-term competitive advantage through negotiation capability and first-party intelligence.

Define programmable deal logic. Establish parameters within which the brand's systems can negotiate: price floors and ceilings, bundle configurations, service-level trade-offs, loyalty incentives, and escalation triggers. Design escalation protocols that define when the system should route to a human negotiator. Build competitive response rules that govern how the system reacts

when an agent presents a competitor's counteroffer. Enable relationship-aware concessions that offer better terms to agents representing high-ULTV customers.

Invest in post-purchase intelligence. Make fulfillment accuracy, delivery speed, return resolution times, and service quality machine-observable — not just internally tracked. Publish structured post-purchase data through agent-accessible feeds. Design connected product experiences that generate continuous usage data. Build feedback loops that convert customer satisfaction signals into structured data the agent's trust model can incorporate. This is the data that feeds agent-trust equity: operational excellence becomes a demand-generation asset.

Secure the first-party relationship. Use the Put-and-Take Method to convert agent-mediated transactions into identified, permissioned customer relationships that feed the brand-owned ecosystem. Design post-purchase touchpoints — connected product activation, warranty registration, personalized follow-up — that work even when the transaction originated through an anonymous agent interaction.

The layered roadmap is calibrated for large enterprises, but the logic applies at every scale. Start with the Foundation Layer: structured data and GEO readiness. The mistake is not starting small. The mistake is starting with the interface — building a chatbot or a branded agent — when the intelligence layer underneath does not yet exist.

7. Forward Look

Three developments will shape the next twelve to eighteen months — and each is an opportunity for the prepared.

Protocol competition will accelerate. UCP, ACP, and MCP differ in governance, scope, and adoption. The winning model will likely be interoperability, not winner-take-all — much as the early web rewarded those who built for multiple browsers.

Late movers won't know what they've lost. The brands building machine legibility now will win decision slots that late movers never realize they lost — because in agentic commerce, the excluded brand is never shown to the consumer. There is no "page two" of results. There is no shelf the consumer walks past. The agent simply never considers you. This is the strategic urgency: the cost of inaction is not visible decline. It is invisible exclusion.

The firms that built Brand Intelligence for humans will discover they built it for machines too. The Command Center that optimizes for human UX can optimize for machine legibility. The smart products that generate behavioral data for personalization generate trust signals for agents. The mobile app that serves as super transmitter sends data to the same Flywheel regardless of which user type initiated the interaction. The investment in the Brand Intelligence architecture was preparation for a world that had not yet arrived. It has now arrived.

Key Concepts at a Glance

Concept	Definition
Machine Legibility	The degree to which a brand's value is structured, accessible, and interpretable by AI agents
Agent Intelligence	The full set of design parameters (Architecture), functional capabilities (Capabilities), and learning inputs (Learning) that determine how an agent shopper decides
Decision-Slot Competition	The struggle to be included in an agent's consideration set before the consumer ever sees the options
Objective-Function Governance	The question of whose objective function controls comparison when the consumer's instructions are ambiguous
Composability Competition	Competition based on the richness of the solution space a brand exposes — more configurable parameters give the agent more ways to find the optimal fit
Deal-Space Competition	Competition based on the breadth of the negotiation parameter space — wider deal architecture gives the agent more dimensions to optimize
Confidence-Model Competition	Competition based on cumulative performance records that compound in the agent's trust model — operational consistency creates incumbent advantage
Agent-Trust Equity	Earned machine trust built through repeated positive outcomes that compound over time
Trust Signal Stack	Five categories of verifiable signals — aggregated outcomes, third-party verification, data quality, system reliability, post-purchase data — that agents use to calibrate trust
Verifiable Value	Differentiation that can be measured, compared, and confirmed through data — as opposed to asserted differentiation
Programmable Deal Logic	Parameters within which the brand's systems can autonomously negotiate: price bounds, bundle configs, concession rules, escalation triggers
Put-and-Take Method	Using external platforms for agent reach (put) while engineering post-purchase conversion to brand-owned relationships (take)

References

1. Klarna, "Klarna Launches Agentic Product Protocol," press release, December 15, 2025. https://www.klarna.com/international/press/klarna-launches-agentic-product-protocol-the-open-standard-that-makes-100m/
2. Google, "New Tech and Tools for Retailers to Succeed in an Agentic Shopping Era," Google Blog, January 11, 2026. https://blog.google/products/ads-commerce/agentic-commerce-ai-tools-protocol-retailers-platforms/
3. Stripe, "Developing an Open Standard for Agentic Commerce," Stripe Blog. https://stripe.com/blog/developing-an-open-standard-for-agentic-commerce
4. PayPal, "PayPal and Perplexity Launch Instant Buy Ahead of Black Friday," press release, November 2025. https://newsroom.paypal-corp.com/2025-11-PayPal-and-Perplexity-Launch-Instant-Buy
5. CNBC, "Amazon Wins Court Order to Block Perplexity's AI Shopping Agent," March 10, 2026. https://www.cnbc.com/2026/03/10/amazon-wins-court-order-to-block-perplexitys-ai-shopping-agent.html
6. Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," press release, August 26, 2025. https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025
7. 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

8. 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

9. 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

10. Walmart, "Walmart and Google Turn AI Discovery Into Effortless Shopping Experiences," January 11, 2026. https://corporate.walmart.com/news/2026/01/11/walmart-and-google-turn-ai-discovery-into-effortless-shopping-experiences

11. 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, 4, 7, 8, 9, 10 | BI-AR-01: GEO Is the New SEO | BI-AR-03: The Agent Divide — Why Agentic Commerce Looks Different in the U.S. and China | BI-AR-04: What Consumers Will — and Won't — Delegate to AI Agents | BI-AR-05: When Agents Negotiate — The New Deal Architecture