BRAND INTELLIGENCE ARTICLE SERIES

How organizations convert everyday interactions into enterprise intelligence.

Interactions → Data → Learning → Intelligence → Advantage

ARTICLE CODE	BI-AR-03
CATEGORY	AI Transformation | Platform Strategy | Global Markets
PILLAR	Competing in the Age of AI
CHAPTER	Ch. 2, 4, 8, 9, 10

The Agent Divide:

Why Agentic Commerce Looks Different in the U.S. and China[1]

Baohong Sun

Summary: During Chinese New Year 2026, Alipay processed 120 million AI-agent transactions in a single week — completed purchases ordered through a chatbot and paid without leaving the conversation. That same week, the U.S. was launching protocols and waitlists. The gap is not about technology. It is about architecture. The United States and China are building agentic commerce through structurally opposite systems — one horizontal, assembling interoperability from fragmented protocols and independent players; the other vertical, activating AI agents inside already-integrated platform ecosystems. This article shows what each system looks and feels like for the consumer, analyzes the distinct infrastructure challenges each creates for brands, and identifies the preparation strategies that hold constant regardless of which system a brand operates in. For executives managing brands across both markets, the core insight is this: the strategic destination is the same — intelligence-driven competitive advantage — but the operational road to get there diverges in ways that demand architecturally aware execution.

Keywords: Agentic Commerce; Horizontal Architecture; Vertical Architecture; Integrated Platform Ecosystem; Agent Divide; Data Sovereignty; Coordination Risk; Concentration Risk; Machine Legibility; Analytical Capability; Flywheel of Intelligence; Put-and-Take Method; Command Center

1. Same Week, Different Worlds

On January 11, 2026, Google unveiled the Universal Commerce Protocol (UCP) at the National Retail Federation conference in New York. The protocol — co-developed with Shopify, Etsy, Wayfair, Target, and Walmart, and endorsed by over twenty additional partners including Mastercard, Visa, Stripe, and American Express — creates a common language for AI agents to discover products, compare options, negotiate terms, and execute transactions across independent commerce platforms.[1] It was designed to sit alongside Anthropic's Model Context Protocol (MCP),

which by December 2025 had reached ninety-seven million monthly SDK downloads and had been donated to the Linux Foundation's Agentic AI Foundation,[2] and Google's own Agent-to-Agent (A2A) protocol, which by mid-2025 counted over 150 supporting organizations.[3]

Four days later, on January 15, Alibaba completed the integration of its Qwen AI model into Taobao, Alipay, Fliggy, and Amap — connecting product search, comparison, payment, travel booking, and navigation inside a single conversational interface.[4] A consumer could ask Qwen to find a winter jacket, compare three options by warmth rating and price, pay through Alipay, and arrange delivery — without leaving the chat window. No protocol handoffs. No interoperability negotiations. The entire commerce stack collapsed into one operator.

Same week. Same ambition. Opposite architectures.

Within a month, the scale difference became unmistakable. During Chinese New Year in February 2026, Alibaba spent 3 billion yuan ($431 million), Tencent spent 1 billion yuan ($144 million), and Baidu spent 500 million yuan ($72 million) subsidizing AI-agent purchases — over $647 million in a single holiday period.[5] Alipay's AI Pay processed 120 million transactions in a single week, surpassing 100 million users.[6] In the U.S. that same month, OpenAI was onboarding its first Etsy merchants for ChatGPT Instant Checkout — and by early March had quietly scaled it back, shifting purchases to retailer apps after only about twelve of Shopify's millions of merchants went live with the feature.[7]

These are not differences of speed. They are differences of structure. And understanding that structural divide is essential for any executive building brand strategy across both markets.

2. Two Shopping Experiences

The fastest way to understand the architectural divide is to watch two consumers shop.

The Chinese Consumer: One Conversation, One Ecosystem

A consumer in Shanghai opens the Qwen app and types: "I need running shoes for a half-marathon next month — comfortable, under 800 yuan." Within seconds, Qwen draws three recommendations from Taobao's product database, each with price, runner reviews, and delivery estimates. The consumer taps on a pair. Qwen asks about size, notes the consumer's purchase history from previous Taobao orders, and pre-fills the selection. The consumer confirms. Alipay processes the payment inside the same chat window. Delivery is dispatched through Cainiao logistics — also an Alibaba company. The consumer never leaves the conversation. No app switching. No login to a separate store. No payment redirect. Total elapsed time: under sixty seconds.

Now consider ByteDance's Doubao agent — with 155 million weekly active users as of early 2026.[8] A consumer asking Doubao for the cheapest price on a specific headphone model triggers something even more striking: the agent opens four competing shopping platforms simultaneously, scans their screens, factors in platform-specific coupons and promotions, compares final

out-of-pocket prices, and presents the cheapest option for one-click purchase. The agent doesn't just search one ecosystem — it shops across ecosystems while remaining inside a single conversational interface.

Meanwhile, Tencent is preparing its own AI agent for WeChat — a project classified internally as top secret — designed to let users hail rides, order food, book restaurants, and shop through conversational prompts within WeChat's mini-program ecosystem, which serves 1.4 billion monthly active users.[9] Gray-box testing is planned for mid-2026.

The pattern across all three Chinese platforms is the same: **the AI agent operates inside an integrated ecosystem where discovery, comparison, payment, fulfillment, and post-purchase service are connected under shared ownership or tight partnership.** The consumer experiences a seamless conversation that produces a completed purchase. The infrastructure complexity is hidden.

The American Consumer: Multiple Handoffs, Multiple Players

A consumer in Chicago asks Google's Gemini: "Find me running shoes for a half-marathon — comfortable, under $120." Gemini, operating through AI Mode in Search, queries multiple merchant backends via the Universal Commerce Protocol. It pulls product data from Nike.com, Zappos, Dick's Sporting Goods, and Hoka's direct store — each of which has separately integrated with UCP's standardized product discovery interface. The results appear in a conversational format with pricing, reviews aggregated from Bazaarvoice, and availability.

The consumer selects a pair. Here is where the experience diverges from China. To complete the purchase, the flow may route to the retailer's own checkout (Google is planning a buy button directly on its AI surfaces via Google Pay, but deeper customization and direct purchasing are still rolling out).[1] If the consumer wants to use an AI agent for the full transaction, they might turn to Perplexity's shopping assistant, which expanded to over 5,000 merchants through PayPal integration[10] — but Amazon just won a federal court order finding that Perplexity's agent accessing Amazon's systems did not constitute authorized access, exposing the legal friction that the horizontal model generates.[11]

The consumer completes the purchase, but it required navigating between AI surfaces, payment systems, and merchant platforms that do not share ownership, do not share data, and interoperate only through voluntarily adopted protocols. The experience works — but the seams show.

What the User Experience Reveals

The difference is not quality. Both consumers got running shoes at a competitive price with reasonable convenience. The difference is *architecture* — and architecture determines who controls what.

In China's **vertical model** (the term "super app" captures the user-facing experience, but the structural reality is a *vertically integrated platform ecosystem* — one operator controlling the AI

interface, product catalog, payment rails, and logistics), the consumer's entire journey occurs within a single operator's system. The platform sees everything: what the consumer searched, what alternatives were considered, which was chosen, how it was paid for, and whether it was returned. The data is unified by default.

In America's **horizontal model**, the consumer's journey crosses independent systems connected by open protocols. Google provides the AI interface. UCP provides the product discovery standard. The retailer owns the catalog. Stripe or PayPal processes the payment. The consumer's data is fragmented across these players by default — unified only if someone builds the infrastructure to aggregate it.

Neither architecture is inherently superior. Each creates distinct advantages and distinct risks for brands.

Why the Gap Is So Wide, So Fast

The scale difference in Section 1 — 120 million AI-agent transactions in China versus protocols and waitlists in the U.S. — is not about technology. It is about infrastructure readiness. China's vertically integrated platforms completed their digital plumbing — mobile payment, integrated commerce, social-to-transaction pathways — before AI agents arrived. Alipay launched in 2004, WeChat Pay in 2013. Chinese consumers entered digital commerce through mobile payment before e-commerce matured, creating an infrastructure base where AI agents deploy into an already-connected system. The result: 120 million AI-agent transactions in a single week, less than three months after Qwen's integration with Taobao.[6]

The U.S. is building the connection layer and deploying AI agents simultaneously. UCP, MCP, and A2A are the connective protocols — horizontal rails that did not exist eighteen months ago. This is structurally more complex: the U.S. must build interoperability while also deploying the agents that depend on it. OpenAI's retreat from native checkout inside ChatGPT — shifting to retailer apps after low merchant adoption — illustrates the challenge.[7]

This does not mean China is "ahead." It means China faces second-stage challenges (platform concentration, sovereignty erosion, algorithmic dependence) while the U.S. faces first-stage challenges (fragmentation, coordination cost, legal uncertainty). Different problems, different timelines. The sections that follow examine what those differences mean for the executives who must build strategy in both systems.

Figure 1: The Agent Divide — Commerce Architecture in Two Markets

Same five layers, opposite ownership structures

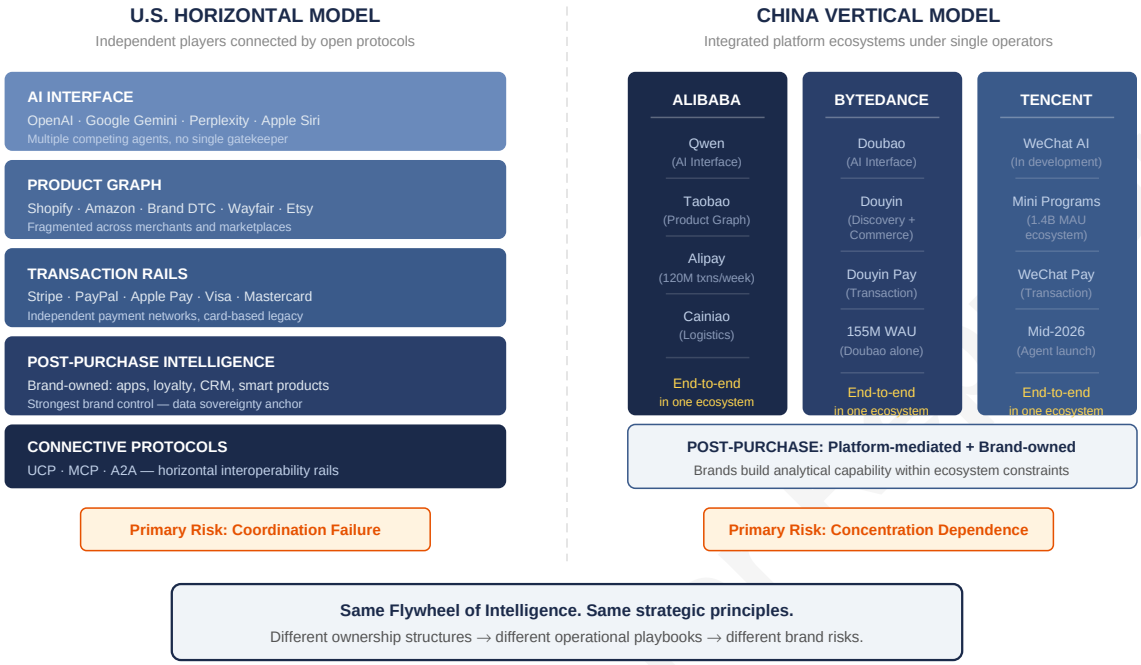

Source: Brand Intelligence (Sun, 2026), Ch. 2, 4, 8, 9. Extended for cross-market agentic commerce comparison.

3. What Differs: Brand Infrastructure in Two Architectures

BI-AR-02 established the infrastructure that brands need for agentic commerce: machine legibility across five agent decisions, a protocol layer serving machine users alongside the human-facing UX, and a Command Center that integrates data into real-time intelligence. That architecture holds in both markets. What changes is the *operating environment* — the constraints, risks, and opportunities that each system imposes on the brand trying to build it.

Machine Legibility: Many Agents vs. One Platform

In the horizontal U.S. system, machine legibility means structuring product data for discovery across multiple independent AI agents — ChatGPT, Gemini, Perplexity, and whatever comes next. Each agent has different data format requirements, different trust calibration mechanisms, and different recommendation algorithms. The brand must maintain legibility across all of them simultaneously, because missing one agent ecosystem means losing the consumer segment that delegates through it. The operational burden scales with the number of agents.

In China's vertical system, machine legibility means structuring data for the platform's native AI — Qwen within Alibaba's ecosystem, Doubao within ByteDance's — in formats that earn preferential recommendation within that specific system. The legibility challenge is narrower but deeper: the brand must understand and optimize for each platform's proprietary algorithm rather than a shared open standard. The formats, API integrations, and performance metrics that earn preferential

recommendation in Taobao differ from those in Douyin or JD.com.

The concept is identical — the brand's value must be structured, accessible, and interpretable by AI agents. The number and type of machine audiences differ.

The Primary Brand Risk: Coordination vs. Concentration

In the horizontal U.S. system, the primary risk is **coordination failure**. A brand must maintain presence across multiple independent agent ecosystems, each with its own protocols, data standards, and commercial terms. Google's UCP, OpenAI's Agentic Commerce Protocol, and Anthropic's MCP represent different integration paths. Missing one means invisible exclusion from every consumer who delegates through it — and the brand may never know it was excluded. The challenge is operational breadth: ensuring the brand is discoverable, transactable, and trustworthy across a fragmented landscape. The Amazon-Perplexity lawsuit reveals an additional dimension: in the horizontal model, the legal boundaries of agent access are still being contested. Brands cannot assume that an open protocol environment means unrestricted agent access to their systems.[11]

In the vertical Chinese system, the primary risk is **concentration dependence**. When a single platform controls the AI interface, product catalog, transaction rails, and recommendation algorithm, the brand's visibility, pricing power, and customer relationship are mediated — and potentially controlled — by one operator. A Taobao flagship store or a WeChat mini-program is simultaneously a brand-building space and a platform-governed environment. The platform can change the algorithm, raise fees, or favor private-label alternatives. The risk is structural vulnerability: too much of the brand's intelligence and customer access concentrated in one counterparty.

Both risks are real. Neither is avoidable. The strategic question is which risk the brand is better equipped to manage — and how to mitigate the one it is most exposed to.

Data Sovereignty: Aggregation vs. Extraction

The strategic imperative of user and data sovereignty takes different operational forms in each market.

In the U.S., sovereignty means building owned infrastructure — brand apps, communities, smart products, loyalty programs — to aggregate intelligence across fragmented sources and reduce dependence on any single platform. The Flywheel of Intelligence (users generate data, data trains algorithms, algorithms improve experiences, better experiences attract more users) spins across multiple independent channels. The Put-and-Take Method operates across these channels: brands broadcast through public domains — social media, AI-discoverable content, borrowed stores — and convert attention into brand-owned ecosystem engagement where the intelligence accumulates. The conversion paths are many, and the orchestration challenge is coordination across independent systems.

In China, the Flywheel and Put-and-Take still apply, but the mechanics shift. A brand operating inside a vertically integrated platform occupies a space that is simultaneously public (platform-mediated, algorithm-ranked) and semi-private (branded, relationship-building). The "put" and "take" may occur within the same ecosystem. The sovereignty goal is not to escape the platform — which is often impractical — but to build **analytical capability** within ecosystem boundaries that the platform cannot replicate or extract.

Analytical capability, in this context, means a specific set of assets: the algorithms trained on the brand's own cross-platform data, the predictive customer models the brand develops from combining platform-mediated interactions with its proprietary intelligence (loyalty data, purchase history across channels, service interactions, smart product usage data), and the ability to generate strategic insights and make decisions that the platform's generic recommendation algorithm cannot. It is what separates a brand that *uses* a platform from a brand that is *dependent on* a platform. A brand with strong analytical capability can operate within Alibaba's ecosystem for reach while running a proprietary intelligence engine that processes the combined data into decisions the platform cannot replicate.

Agent-Trust Equity: Distributed vs. Concentrated

BI-AR-02 introduced agent-trust equity — the cumulative performance record that determines whether an AI agent will recommend a brand in future transactions. How that trust forms differs by architecture.

In the horizontal U.S. model, agent-trust equity builds across multiple independent agents. A brand must earn trust from ChatGPT, Gemini, and Perplexity separately — each evaluates the brand's machine legibility, fulfillment reliability, and outcome data through different mechanisms. Trust is distributed and non-transferable. A strong reputation with one agent does not automatically transfer to another.

In China's vertical model, agent-trust equity concentrates within platform ecosystems. A brand's Taobao trust score, Douyin reputation metrics, and JD.com performance data are platform-specific and largely non-portable. The brand must build trust within each ecosystem separately — but within each, the trust data is richer and more unified because the platform sees the entire transaction lifecycle.

4. What Holds Constant: Preparation Strategies for Both Systems

The architectural differences are real, but they sit on top of strategic principles that do not change regardless of ecosystem structure. Executives who grasp these invariants can build one strategic logic with two operational playbooks.

Intelligence Accumulation Is the Compounding Asset

The Flywheel of Intelligence operates identically in both systems. Whether it spins inside an integrated platform or across a fragmented protocol network, the brands that build proprietary behavioral data and train proprietary algorithms accumulate advantage that compounds over time. The Flywheel is architecture-neutral: what matters is that the brand owns the intelligence it generates, not which pipes the data flows through.

Post-Purchase Intelligence Is the Decisive Competitive Asset

In both systems, the brands that capture outcome data — fulfillment accuracy, return rates, satisfaction signals, usage patterns — compound their advantage through the Trust Signal Stack that BI-AR-02 described. A product that generates continuous usage data (driving patterns, charging behavior, app interaction, subscription renewal) feeds the Flywheel regardless of which commerce ecosystem surrounds it. Post-purchase intelligence is the hardest asset for competitors to replicate and the most valuable for earning agent-trust equity — in both architectures.

The Command Center Architecture Applies Everywhere

The operational brain that integrates data into real-time algorithmic marketing intelligence — the Customer Data Platform unifying first-party data, the Algorithmic Deployment Center operationalizing predictive models — is the same strategic answer in both systems. What differs is whether it aggregates across independent sources (U.S.) or generates proprietary intelligence within ecosystem constraints (China). The function is identical; the data plumbing adapts.

Consumer Delegation Follows Behavioral Logic, Not Infrastructure Logic

The categories where consumers will delegate purchase decisions to AI agents — commodity replenishment, price-sensitive comparison shopping, routine reorders — follow from the structure of the decision, not the structure of the platform. High-emotional-investment categories (luxury, personal health, education) resist delegation in both markets. Cultural overlays modify speed but not structure. This means a brand's assessment of which decisions agents will handle is portable across markets.

Machine Legibility Is Non-Negotiable in Both Systems

Whatever the architecture, a brand that is not structured, accessible, and interpretable by AI agents does not exist in the agent's consideration set. The five legibility requirements — findability, comparability, modularity, actionability, and verifiability — hold universally. The number and type of machine audiences are architecturally contingent; the requirement to be legible is not.

5. Executing Across Both Systems

For executives managing brands in both the U.S. and China, the operational challenge is translating universal principles into architecturally appropriate execution. Four imperatives apply across both markets — each with system-specific requirements.

Build the intelligence layer, not the interface. In the U.S., this means investing in the Command Center — a Customer Data Platform that unifies first-party data across fragmented sources, structured data feeds that ensure machine legibility across multiple AI agent ecosystems, and the analytical capability to act on that data independently. In China, it means building proprietary analytical capability within platform boundaries — the algorithms, customer models, and predictive intelligence that the brand owns even though the raw data flows through platform-mediated channels.

Execute Put-and-Take with architectural awareness. In the horizontal U.S. system, "put" happens across fragmented channels — social media, AI-discoverable content, borrowed stores — and "take" routes through brand-owned digital properties. In China's vertical system, "put" and "take" may occur within the same ecosystem, but the brand must engineer the conversion from platform-mediated interaction to brand-owned relationship. The mechanism differs; the strategic logic is the same.

Design for machine legibility in the right number of systems. In the U.S., machine legibility must be maintained across multiple independent agent systems — a coordination challenge. In China, machine legibility is platform-specific — a depth challenge. Global brands need both capabilities, which argues for a centralized machine-legibility function that produces architecturally adapted outputs for each market.

Figure 2: Four Imperatives for Executing Across Both Systems

Same principles, architecturally specific execution

IMPERATIVE	U.S. HORIZONTAL	CHINA VERTICAL
1. Build the Intelligence Layer Command Center (Ch. 9) as competitive anchor	CDP unifying first-party data across fragmented sources; machine legibility across multiple independent agents	Proprietary algorithms within integrated ecosystem; analytical capability the platform cannot replicate
2. Execute Put-and-Take Omni-domain strategy (Ch. 10)	"Put" across many channels; "Take" through brand-owned digital properties. Many conversion paths to orchestrate.	"Put" and "Take" may occur within same ecosystem. Engineer conversion from platform to brand relationship.
3. Design for Machine Legibility Structured data for AI agent audiences	Maintain legibility across multiple independent agents: ChatGPT, Gemini, Perplexity. Coordination challenge.	Platform-specific legibility: Taobao ≠ Douyin ≠ JD.com. Fewer agents, deeper rules. Depth challenge.
4. Protect Data Sovereignty User & Data Sovereignty (Ch. 1, 4)	Own the infrastructure that aggregates intelligence. Risk: data scattered across platforms brand cannot unify.	Build analytical capability platform cannot replicate. Risk: single platform mediates too much of brand intelligence.

Source: Brand Intelligence (Sun, 2026), Ch. 1, 4, 8, 9, 10. Extended for dual-system agentic commerce strategy.

Protect data sovereignty through architecturally appropriate strategies. In the U.S., sovereignty means owning the data infrastructure that aggregates intelligence across fragmented

sources. The risk is fragmentation: data scattered across platforms that the brand cannot unify. In China, sovereignty means building analytical capability that the platform cannot replicate or extract. The risk is concentration: a single platform mediating so much of the brand's intelligence that the brand becomes operationally dependent. In both cases, the Command Center is the strategic answer — the difference is whether it aggregates across independent sources or generates proprietary intelligence within ecosystem constraints.

6. Forward Look: Convergence or Divergence?

Three developments will determine whether the horizontal and vertical models move toward each other or further apart.

The legal boundaries of agent access are being drawn now. The Amazon-Perplexity case is the first major legal test of agentic commerce in the U.S. — and the initial ruling (that AI agents accessing a platform on a user's behalf did not constitute authorized access) could reshape how the horizontal model operates. If platforms can legally block external agents, the horizontal model develops gatekeepers that begin to resemble vertical control. The 9th Circuit appeal, still pending, will set precedent that every brand strategist should track.[11]

Protocol adoption in China could fragment the vertical stack. If Alibaba, Tencent, or ByteDance adopt open interoperability protocols — MCP, UCP, or Chinese equivalents — the vertical stack may begin to open. Alibaba's March 2026 launch of the Wukong enterprise AI platform, with plans for Slack and Teams integration, signals at least selective openness to cross-ecosystem connectivity.[12] But the core consumer-facing ecosystems remain closed. The question is whether competitive pressure or regulatory mandate forces interoperability.

Regulatory trajectories could push convergence. China's Personal Information Protection Law (PIPL) and evolving U.S. state-level privacy legislation are both tightening data governance, but through different mechanisms. If regulatory convergence occurs, the data sovereignty architectures may align. If they diverge further, global brands will need genuinely separate data strategies for each market.

The strategic logic holds regardless of architectural convergence: the brands that accumulate intelligence faster than the market — through the Flywheel, the Command Center, and the Put-and-Take Method — will define the competitive landscape in both systems. The architecture is contingent. The intelligence is permanent.

Key Concepts at a Glance

Concept	Definition
Horizontal Architecture	Commerce stack with independent players (AI interface, product catalog, payments, logistics) connected by open protocols (UCP, MCP, A2A) — the U.S. model
Vertical Architecture	Commerce stack with multiple layers controlled by a single integrated platform ecosystem — the Chinese model (commonly called "super apps")
Coordination Risk	Primary brand risk in horizontal systems: maintaining machine legibility and presence across multiple independent agent ecosystems
Concentration Risk	Primary brand risk in vertical systems: dependence on a single platform that controls multiple layers of the commerce stack
Analytical Capability	The brand's proprietary algorithms, predictive models, and cross-platform intelligence that enable strategic decisions the platform's generic algorithm cannot replicate
Machine Legibility	The degree to which a brand's value is structured, accessible, and interpretable by AI agents — required in both architectures
Agent-Trust Equity	Cumulative performance record that determines agent recommendation probability — builds differently (distributed vs. concentrated) by architecture
Flywheel of Intelligence	Self-reinforcing loop: users → data → algorithms → better experiences → more users — operates identically in both systems
Put-and-Take Method	Broadcasting through public channels ("put") and converting attention into brand-owned ecosystem engagement ("take") — same logic, different mechanics by architecture
Command Center	Operational brain integrating customer data and algorithmic deployment for real-time marketing intelligence — the strategic answer in both systems

References

1. Google, "New Tech and Tools for Retailers to Succeed in an Agentic Shopping Era," Google Blog, January 11, 2026. Partners include Shopify, Etsy, Wayfair, Target, Walmart, Mastercard, Visa, Stripe, American Express. UCP documentation: developers.google.com/merchant/ucp
2. Anthropic (2025). Model Context Protocol donated to the Agentic AI Foundation under Linux Foundation governance, December 2025. By December 2025, MCP had reached 97 million monthly SDK downloads across Python and TypeScript.
3. Google, "A2A: A new era of agent interoperability," Google Developers Blog, April 9, 2025. By mid-2025, A2A counted 150+ supporting organizations.
4. Caixin Global, "Alibaba Integrates AI Chatbot With Taobao, Alipay," January 16, 2026. Qwen connected to Taobao, Alipay, Fliggy, and Amap on January 15, 2026.
5. AI-News, "China's hyperscalers bet billions on agentic AI as commerce becomes the new battleground," February 2026. Alibaba 3B yuan, Tencent 1B yuan, Baidu 500M yuan. Rest of World, "Free boba and red envelopes: China's AI giants launch marketing blitz for Lunar New Year," February 2026.
6. BusinessWire, "Alipay AI Payment Exceeds 120 Million Transactions in One Week as Agentic Commerce Accelerates in China," February 13, 2026. AI Pay surpassed 100 million users by February 24, 2026.
7. OpenAI, "Buy it in ChatGPT: Instant Checkout and the Agentic Commerce Protocol," 2026. 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.

8. ByteDance, "Doubao 2.0 AI model for 'agent era,'" February 14, 2026. QuestMobile data: 155 million weekly active users. Doubao leads all AI chatbot apps in China, with DeepSeek second at 81.6 million.
9. TechNode, "Tencent is said to be developing a top-secret AI agent project for WeChat," March 11, 2026. WeChat's mini-program ecosystem serves 1.4 billion monthly active users. Nikkei Asia, "Tencent developing AI agent for WeChat as OpenClaw fever grips China," March 2026.
10. PayPal, "PayPal and Perplexity Launch Instant Buy Ahead of Black Friday," press release, November 2025.
11. CNBC, "Amazon Wins Court Order to Block Perplexity's AI Shopping Agent," March 10, 2026. Bloomberg, "Perplexity AI Shopping Bots Can Stay on Amazon for Now, Court Rules," March 17, 2026. The 9th Circuit granted a stay of the lower court order pending appeal.
12. CNBC, "Alibaba launches agentic AI tool for businesses with Slack, Teams integration plans," March 17, 2026. Wukong enterprise AI platform for business agent management.
13. 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, 4, 8, 9, 10 | BI-AR-01: GEO Is the New SEO | BI-AR-02: The Agent Shopper — Brand Strategy for the Agentic Commerce Era