When the buyer is an algorithm, the rules of the game change

How agentic commerce changes buying behaviour, and how brands, retailers and service businesses win

For decades, brands won by following a combination of Byron Sharp's laws: be easy to find and buy (physical availability) and be salient, distinctive, top-of-mind (mental availability). Both assume a human makes the call. This paper asks what rules will apply when an agent chooses instead. Does mental availability still matter to a machine that is "immune to nostalgia", or does it collapse into structured-data availability, feed quality and machine-readable trust? Or does it matter even more, because you want your shopper steering the agent towards your brand in the first place?

Executive Summary

For nearly two decades, Byron Sharp's How Brands Grow has functioned as marketing's closest thing to settled science: brands grow by building physical availability (being easy to find and buy wherever the category is bought) and mental availability (being easy to bring to mind through distinctive assets and broad, undifferentiated reach). Both laws rest on an assumption so basic it was never stated out loud: a human being stands at the point of choice, scanning a shelf or a search results page, deciding what to buy. That assumption is now failing, category by category, as a growing share of purchase decisions are made not by the shopper but by an AI agent acting on the shopper's behalf.

The shift is measurable today. ChatGPT alone reached roughly 900 million weekly users by early 2026, up from around 100 million eighteen months earlier. The share of consumers using generative AI somewhere in their shopping journey rose from 38% in 2024 to 51% in 2025, with four in five planning to do so in 2026. AI-referred traffic to retail sites grew more than 300% over the course of 2025, while roughly two-thirds of Google searches now end without a single click to any website. McKinsey puts the addressable prize at up to $1 trillion of orchestrated US retail revenue and $3 to 5 trillion globally by 2030; Bain, Morgan Stanley, BCG and Gartner each arrive at different numbers through different methods, but all point in the same direction and at a similar order of magnitude.

This paper's central claim is that the shift rewrites the terms on which physical and mental availability are won. Physical availability survives but relocates, from pallet and shelf to product feed, API and payment protocol. Mental availability bifurcates: it continues to matter for the human who sets the goal and decides whether to trust a brand enough to delegate a decision to

an agent in the first place, but for the agent that actually executes the in-basket comparison, salience built on emotional memory counts for almost nothing. What counts there is structured, verifiable, machine-readable data. As one analysis bluntly puts it, AI agents are immune to nostalgia.

The ground beneath this shift is also still moving. In March 2026, OpenAI scaled back ChatGPT's Instant Checkout after its 4% transaction fee stalled merchant adoption, while Perplexity's free "Buy with Pro", funded by subscriptions rather than transaction fees, gained share. Walmart is building an open, multi-platform agent strategy; Amazon is building a closed one. Brands, retailers and service businesses are being asked to place strategic bets on infrastructure that is itself still contested. What follows lays out the evidence for this shift, its implications by business type and function, and a 30/60/90-day path for getting ahead of it rather than reacting to it.

1. The Rise of Agentic Commerce

For twenty-five years, e-commerce infrastructure was built for a single kind of user: a human, scrolling through filters, comparing prices across browser tabs, abandoning carts when the process became annoying. "Now AI agents translate intent, weigh constraints, and make commitments on behalf of users," as the team behind The Fold, a vendor building reasoning infrastructure specifically for agent-mediated commerce, frames the shift. Agentic commerce, in this sense, is commerce in which the agent itself is the customer of record for a given transaction: the entity that reads the product page, weighs the alternatives, and clicks buy, even though the money and the eventual unboxing both belong to a person who may never have looked at the page at all.

The rails arrived faster than anyone expected

The infrastructure arrived from every direction at once. OpenAI and Stripe co-developed the Agentic Commerce Protocol and launched Instant Checkout inside ChatGPT in September 2025, with Stripe noting that any merchant already on its platform could enable agentic payments "in as little as one line of code." Visa followed with Intelligent Commerce, embedding payment credentials, authentication and fraud controls directly into agent-initiated transactions, then consolidated four competing agent protocols, including its own Trusted Agent Protocol built in partnership with Akamai's bot- and behaviour-detection infrastructure, into a single integration called Intelligent Commerce Connect. Mastercard answered within weeks with Agent Pay, built on tokenised credentials and its own Payment Passkeys. Google released its Agent Payments Protocol, and in June 2026 Adyen launched Adyen Agentic, which it calls "the universal translator for the next era of commerce," built around an Agentic Feed, an Agentic Cart and Agentic Payments that together translate a merchant's catalogue into a form an agent can actually reason over. Microsoft took a different entry point, building Copilot Checkout, which lets a shopper complete a purchase inside the Copilot app itself without ever leaving the conversation, alongside Brand Agents that put a merchant's own AI assistant directly onto its product pages. For a brand or retailer, this is the new physical availability: a clean, structured,

API-reachable line item an agent can discover, evaluate and check out against in one call, rather than a place on a shelf.

Big, and disputed, numbers

The size of the prize is large enough that almost every major analyst house has put a number on it, and the numbers disagree more than they agree. McKinsey's estimate runs to $1 trillion of orchestrated US retail revenue and $3 to 5 trillion worldwide by 2030. Bain is more conservative on the US figure, at $300 to 500 billion or 15 to 25% of e-commerce, while Morgan Stanley lands lower still at $190 to 385 billion, adding the detail that agentic spending concentrates first in groceries and consumer packaged goods, the categories where loyalty was always thinnest. BCG puts the influenced-spending figure at $1.3 trillion, and Gartner, looking past consumer retail entirely, forecasts that AI agents will intermediate more than $15 trillion of B2B purchasing by 2028. The forecasts disagree on precision but agree on direction: every serious analyst house treats agent-mediated purchasing as a multi-trillion-dollar inevitability rather than a niche behaviour to monitor from a distance.

The rails are a business-model fight, not neutral plumbing

The clearest evidence that the infrastructure remains unsettled is a business-model fight still being lost and won in real time. OpenAI built its Agentic Commerce Protocol around a 4% transaction fee charged to merchants; adoption stalled, and in March 2026 OpenAI scaled back Instant Checkout rather than continue pushing a fee structure the market would not absorb. Perplexity took the opposite bet. Its "Buy with Pro" feature, built with PayPal, charges merchants nothing at all, recovering the cost instead through Pro subscription revenue.

> *"Perplexity charges merchants nothing."*
>
> Source: Stellagent, 2026

A parallel and equally unresolved bet is being placed at the retailer level. Walmart has chosen openness: its own "Sparky" agent sits alongside live integrations with ChatGPT, Gemini and the Universal Commerce Protocol, trading some control for the broadest possible reach across whichever agent a shopper happens to be using. Amazon has chosen control: its agent strategy, led by "Buy for Me," keeps discovery, comparison and checkout inside its own environment from end to end.

> *"Walmart plugs into external AI ecosystems (ChatGPT, Gemini, Google's Universal Commerce Protocol) while also building internal agents. [Amazon] focuses on keeping transactions and AI interactions inside its own environment, from search to checkout."*
>
> Source: ALM Corp, 2026

The Walmart story is worth telling in more detail, because it is also a cautionary one. Walmart ran roughly 200,000 products through ChatGPT's in-chat purchasing flow, and those purchases converted at only a third of the rate Walmart sees on its own site, undone largely by wrong

delivery estimates and incorrect shipping costs surfaced inside the chat. The lesson is that agentic checkout fails immediately and visibly when the underlying data feed is wrong, with no salesperson on hand to smooth over a bad number. A similar dynamic played out in China, where ByteDance's agent-enabled phone sold its first 30,000 units within hours before WeChat, Taobao and Alipay each moved within days to restrict the agent's access to their platforms, citing security but, as one analysis of the episode notes, plainly protecting their own margins from a competitor's agent doing the shopping on their turf.

Neither approach, openness or control, has yet proven decisively superior, and any brand or retailer building its own agent strategy today is making a bet on which philosophy will define the next decade of retail, whether it acknowledges that or not.

2. What Is Changing Today

This is not a forecast exercise. Adoption of the major conversational AI platforms has compounded rather than plateaued: ChatGPT, Gemini, Perplexity and Claude have each posted month-on-month user growth through the first half of 2026, and AI-assisted search, AI-generated recommendations and AI-mediated shopping have moved from novelty to mainstream within roughly two years. Industry commentary from events such as Shoptalk 2026 increasingly treats agentic shopping and outcome-driven AI as the operating reality retailers must already plan around, rather than as emerging trends to watch from a distance.

The flip side of agent-mediated discovery is that the channel it replaces is shrinking. A clear majority of Google searches now end without a single click to any website, and publisher search traffic fell by roughly a third globally in the year to November 2025. Any business whose acquisition strategy still assumes a human will click a blue link, browse a results page and form an impression along the way is building on ground that is eroding in real time. The same pattern is appearing on the enterprise side, where procurement agents, customer-service agents and personal AI assistants follow the same delegation curve as consumer shopping, which is why Gartner's $15 trillion B2B forecast deserves equal weight to the consumer figures that dominate press coverage. The three trend lines that matter most, adoption growth, the rising share of agent-influenced shopping journeys, and the shift in where discovery happens, are tracked on the live dashboard accompanying this paper.

3. How Buying Behaviour Is Changing

Buying journeys are moving along a gradient rather than making a single jump from human-only to fully autonomous: human-only, to human-plus-AI, to AI-assisted, to agent-driven, to autonomous purchasing. Different categories sit at very different points on that gradient at the same moment in time, and delegation, where it happens, is partial and trust-gated rather than total. Bain's own consumer research finds that around half of shoppers remain cautious about letting an agent handle a purchase autonomously from start to finish, even as 30 to 45% of US consumers already use generative AI for product research and comparison. Other surveys put

the research-and-comparison figure even higher: by one estimate, 73% of consumers already use AI somewhere in their shopping journey, most commonly for product ideas, review summarisation and price comparison, with 70% saying they are at least somewhat comfortable with an agent eventually making the purchase itself (Ekamoira, 2026).

> *"73% of consumers are already using AI in their shopping journey."*
>
> Source: Ekamoira, 2026

The order in which categories move along this gradient is not random, and it is unflattering to the categories that built their value on emotional differentiation. McKinsey's own automation curve, mapping how much of a purchase decision consumers are willing to hand to an agent category by category, shows delegation arriving first and fastest in low-emotion, high-repeat purchases: replenishment, commodities, groceries and CPG, exactly the categories Morgan Stanley identifies as the largest near-term growth driver for agentic spending. These are also, not coincidentally, the categories where brand attachment was always thinnest; nobody forms a deep emotional bond with a brand of paper towels, so there is little loyalty for an agent to override.

Delegation arrives last in identity-laden, desire-driven categories: luxury, fashion, anything purchased partly for what it signals about the buyer rather than what it does. McKinsey's own research into luxury in the agentic age finds that these categories, built on aura and human-centred desire, are the hardest for an attribute-matching agent to flatten, because the thing being bought is not really the product specification at all. Delegation is arriving fastest exactly where emotional attachment to a brand always mattered least, and arriving slowest exactly where it always mattered most.

4. How Winners Will Win

The single most consequential behavioural claim emerging from this body of evidence is that AI agents do not carry the cognitive biases that mental availability has always exploited. Human memory is associative, emotional and lazy, reaching for the brand it recognises rather than the brand that is objectively best, and that laziness is the entire mechanism by which decades of brand advertising has worked. An agent has no such laziness to exploit.

> *"AI agents are immune to nostalgia."*
>
> Source: Bazaarvoice, 2025

This claim has now been tested directly, not only argued in theory. Researchers Jafar Sabbah of Bayes Business School and Oguz A. Acar of King's College London ran thousands of simulated shopping rounds across four leading AI models and four product categories, pitting the classic toolkit of e-commerce persuasion, scarcity messages, countdown timers, vouchers, bundling and strike-through pricing, against agent decision-making. The tactics built for human psychology did not reliably work, and several backfired outright depending on the model and

the category, with the more advanced reasoning models often the most openly sceptical of overt persuasion. Only authentic star ratings consistently increased an agent's likelihood of choosing a product; a higher price reliably decreased it, exactly as a rational comparison would predict. The authors' conclusion for marketers is to treat each AI model as a distinct audience with its own quirks, double down on fundamentals such as competitive pricing and credible reviews, and invest in testing infrastructure that keeps measuring how different agents respond as the underlying models keep changing.

> *"In thousands of simulated shopping rounds... only star ratings consistently increased choice in the expected direction, while price reliably decreased it."*
>
> Source: Harvard Business Review, "Research: Traditional Marketing Doesn't Work on AI Shopping Agents," May 2026

Byron Sharp himself has weighed in on this question, framing AI agents as simply the latest in a long line of distribution intermediaries brands have had to win over, no more existentially novel, in his telling, than the bartenders, hardware sales people, doctors and mortgage brokers who have long stood between a brand and its eventual buyer. One LinkedIn commentary on the same research distilled the implication into a single line worth keeping in view through the rest of this paper:

> *"The question isn't whether your brand can persuade an algorithm. It's whether it's strong enough not to need to."*
>
> Source: commentary on HBR research, June 2026

A separate framework for what it calls the "dual-interface brand" makes a related point more simply: agents weigh a prompt's specific parameters rather than broad, mass-appeal marketing claims, so the data behind a product matters more than the story told about it. Taken together, this paper reads the evidence as a bifurcation rather than a disappearance. Mental availability continues to matter, perhaps even more than before, for the human who sets the goal and decides whether to trust a brand enough to delegate the decision to an agent at all; brand reputation is what earns the right to be considered before the agent ever opens a comparison. For the agent actually executing the in-basket comparison once delegation has happened, salience built on emotional memory is close to worthless. What the agent rewards instead is what Accenture's own framing calls being "machine-legible and trusted": clean, structured, verifiable data and signals an algorithm can parse with confidence, rather than a feeling a human can recall.

This claim is increasingly measured rather than directional. One analysis of Google's evolving Merchant Center requirements finds that stores achieving what the industry now calls a "Golden Record," 99.9% completion across product attributes, see three to four times higher visibility in AI shopping recommendations than stores with sparse data.

> *"Stores with 99.9% attribute completion (what the industry calls a "Golden Record") are seeing 3-4x higher visibility in AI recommendations compared to stores with sparse data."*
>
> Source: eFulfillment Service, 2026

Google has responded to this dynamic by adding dozens of new attributes to Merchant Center built specifically for conversational commerce: answers to commonly asked product questions, compatible accessories, suitable substitutes, extending the feed from a price-and-availability list into something closer to the actual language an agent reasons in. Brand-side infrastructure is catching up to the same insight. Akamai's AI Brand Presence product, for instance, exists to detect AI agents at the network edge and serve them optimised, structured content distinct from what a human visitor sees, on the reasoning, stated plainly in its own marketing, that "AI systems don't browse like humans. They parse, summarize, and generate responses, which requires structured, machine-readable content." The product also tracks a brand's citation rate across ChatGPT, Claude, Gemini and Perplexity, treating share of voice inside an agent's answer as a metric every bit as trackable, and every bit as contestable, as share of shelf once was.

Old rule (human shopper)	New rule (agent shopper)
Win the shelf and retail-media placement	Win the answer / be in the agent's consideration set
Distinctive assets for memory	Structured attributes for machine matching
Brand salience and emotional reach	Trustworthy reviews and verifiable claims
Be easy to buy (distribution)	Be transactable by an agent (feeds, APIs, protocols)

This advantage is not permanent or platform-neutral, which is the second half of the argument and the one most often missed. The infrastructure a brand optimises for today may not be the infrastructure that wins tomorrow. The OpenAI/Perplexity fee divergence and the Walmart/Amazon openness divergence, both described above, mean that a brand betting everything on one agent ecosystem's rules is making a strategic bet on an unresolved war, not building on settled ground.

Five moves

Taken together, the evidence points toward five concrete moves.

- Serve two customers: the human who decides whether to delegate, and the agent that executes.
- Treat your feed as your flagship store. Structured, accurate, attribute-level data is the new distribution, so aim for a Golden Record rather than a minimum-viable listing.
- Build machine-legible trust through reviews, verifiable claims and protocol-level identity. These become brand equity.
- Win Generative Engine Optimisation by making your brand discoverable and citable inside the agent's answer.

- Hedge across open and closed ecosystems. The OpenAI/Perplexity fee war and the Walmart/Amazon platform split are not resolved, so a strategy built for only one outcome is a bet rather than a plan.

5. Implications for Product Businesses

FMCG (e.g. P&G, Unilever)

For FMCG businesses, the most consequential battle is moving out of the advertising budget and into the recommendation algorithm and the replenishment loop. Shelf visibility, the physical-availability metric that has governed grocery retail for a century, is becoming feed visibility, a function of how completely and accurately a product's attributes are represented in the data an agent reads rather than how prominently it is displayed at eye level. Subscription and automatic replenishment, already meaningful before agentic commerce, are becoming the default basket once an agent is able to reorder without being asked, and Gartner's forecast that 60% of brands will use agentic AI for one-to-one customer interaction by 2028 suggests this is rapidly becoming table stakes rather than a differentiator. Brand power, in this world, is increasingly expressed as data power: the completeness of a feed, the credibility of a review corpus, the verifiability of a claim. Microsoft's Brand Agents, now live for merchants on Shopify, and its Copilot Checkout, which lets a shopper complete a purchase entirely inside the Copilot app without redirecting to external websites, are an early, concrete illustration of what this looks like in practice: the brand itself becomes a conversational surface, not merely a listing the agent reads.

The scale of this shift is already being budgeted for inside the FMCG companies themselves. Mondelez now projects that 20 to 30% of its product purchases will happen via AI agents within one to two years, and has created a dedicated global lead role for emerging commerce platforms to build the playbook for it. Google, for its part, has moved to formalise the new battleground directly inside Merchant Center, launching Conversational Attributes (natural-language product data tuned for how people actually phrase requests to an agent), AI Performance Insights (a share-of-voice metric comparing a brand's visibility on AI surfaces against named competitors), and Universal Cart (a checkout that works across retailers via Search and Gemini, already live with partners including Nike, Sephora, Target and Walmart). One CPG-side commentator on this shift put the change in distinctly personal terms:

> *"I barely walk into a supermarket anymore. My agent does it."*
>
> Source: LinkedIn commentary, June 2026

Durable Goods (automotive, electronics, appliances)

For durable goods businesses, the comparison itself is what is being automated. Configuration and comparison agents are increasingly doing the spec-by-spec evaluation that a showroom visit or a glossy brochure used to win through presentation and salesmanship, and the product that wins is the one whose specifications are clearest and most completely represented rather than

the one with the best-looking advertisement. Ecosystem lock-in becomes a more important determinant of what an agent even considers, since an agent reasoning about compatibility will naturally favour products that integrate cleanly with what a customer already owns, and service revenue becomes the durable margin once the hardware purchase decision has been substantially commoditised by comparison agents that treat every competing product as a comparable set of numbers.

6. Implications for Service Businesses

Service businesses face a version of the same dynamic, complicated by the fact that the product being compared is itself often a promise rather than a physical object.

Travel

AI trip planners and dynamically generated packages are shifting real booking volume toward agent-to-agent transactions, in which a consumer's personal agent negotiates directly with a supplier's booking system. The winning supplier in that negotiation is the one whose availability, pricing and terms are most legibly exposed via API, not the one with the most memorable advertising campaign.

Insurance

AI brokers and product-comparison agents are doing to policy shopping what comparison sites already began a decade ago, only faster and with less human patience for ambiguity, while claims automation is reshaping the cost side of the business at the same time the front end is being reshaped.

Banking & Wealth

Personal finance agents are beginning to drive automated switching and even automated advice, and customer loyalty, historically one of the stickiest assets in financial services, erodes meaningfully once an agent is continuously and tirelessly shopping the market on a customer's behalf rather than waiting for the customer to get annoyed enough to switch manually.

Forrester's own assessment of agentic payments readiness in B2C commerce is a useful corrective here. The infrastructure described above is real, but it is not yet uniformly mature across the discovery-to-checkout flow, and service businesses with longer, more relationship-dependent sales cycles have more runway than the headline statistics might suggest. That runway should be used to prepare deliberately rather than treated as evidence the shift can be ignored.

7. Functional Implications

Agentic commerce touches strategy, technology, finance, marketing and commercial operations simultaneously, and each function has a different lever to pull. It is not a project that can be handed to a single function and monitored from a distance.

CEO	**Own the strategic question and exposure; decide where to own the relationship versus meet the third-party agent.**
CIO	APIs, MCP, agent identity, security, agent-readiness, architecture.
CFO	New economics, shifting revenue pools, agent-era metrics, investment priorities.
CMO	Generative Engine Optimisation, brand discoverability, zero-click marketing, content architecture.
CCO	Commercial model, pricing, channel conflict, the agent-era sales model.

8. Risk, Regulation & Trust

Trust is the gate through which all of the preceding analysis must pass, and the regulatory architecture meant to provide that trust was built for an entirely different kind of transaction. As of early 2026, no jurisdiction has enacted regulation specifically addressing agentic commerce, and the patchwork of existing rules, the EU AI Act, PSD3, GDPR and the Consumer Rights Directive in Europe alone, overlap without resolving the most basic question the new model raises: when an AI agent makes an unauthorised or harmful purchase, who is actually liable?

> *"When an AI agent makes an unauthorised or harmful purchase, who is liable: the consumer, the AI provider, the merchant or the platform?"*
>
> Source: European Business Magazine, 2025

Legal analysts surveying the contract-formation question are no closer to consensus, since an agent's purchase raises genuine doubt about whether a binding contract has even been formed in the way consumer-protection law assumes. The manipulation risk runs in both directions at once, which is what makes it genuinely novel rather than a simple extension of existing concerns. The same agent that is immune to a brand's nostalgia marketing can, in principle, be gamed through the data it is fed rather than the emotions it is appealed to: one survey of retailer risk puts the AI agent manipulation vulnerability rate at roughly 25%, while financial institutions are bracing for a corresponding rise in agent-linked fraud. Accenture's own research finds 78% of financial institutions expect fraud connected to AI agents to increase, and 87% of CTOs and payments leaders believe trust, not technology, will be the binding constraint on how fast agentic payments scale. Persuasion, the soft, manipulative core that mental availability has always partly relied on, reappears in an agentic world as an attack surface to defend rather than a lever brands can pull directly. The platform, protocol or brand that manages to make agentic delegation feel genuinely safe to the human granting it will earn spending the rest of the market is still fighting for.

9. 30 / 60 / 90 Day Playbooks

None of the preceding analysis is useful without a path to act on it, and the path does not need to start with a multi-year transformation programme.

First 30 days: Understand. Educate leadership, map buying journeys, assess exposure, surface assumptions. Deliverable: an Agentic Commerce Assessment.

Days 31 to 60: Experiment. Test AI discovery, build APIs, create GEO content, explore use cases. Deliverable: pilot initiatives.

Days 61 to 90: Operationalise. Define the roadmap, create governance, allocate budget, build capabilities. Deliverable: a 12-month roadmap.

Sources

This paper is grounded in the curated source library (49 sources and growing), spanning consumer behaviour, brand strategy, retailer impact, technology and payment rails, and regulation. It draws on primary research from McKinsey, Bain, BCG, Morgan Stanley, Gartner, Forrester, Accenture, PwC and Harvard Business Review; company sources from OpenAI, Stripe, Visa, Mastercard, Adyen, Akamai, Microsoft, Google, Shopify and The Fold; and reporting, analysis and commentary from Bazaarvoice, Sia Partners, eMarketer, Stellagent, ALM Corp, eFulfillment Service, Howard Yu and Byron Sharp, among others. Figures and quotes are to be verified at review before publication. The full, continuously updated library is available on the Source Library page of this site.

This paper is developed by IntotheNXT, the consultancy helping companies to be more entrepreneurial and innovative by helping leaders to understand technology trends and shape companies where humans bring their best.

Amsterdam, June 2026