ACTIVANT RESEARCH

Agentic Commerce

Market to people, Sell to Software

July 2026

AUTHORS:

Steve Sarracino

Founder & Partner

Torben Wiesbach

Senior Research Analyst

Scott Watson

Vice President, Research

About Activant

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at https://activantcapital.com/.

In the [first article](#) in this three-part series we argued that agentic commerce is a stack of three tiers, that agents will dominate some consumer categories and remain largely irrelevant in others, and that the line between those outcomes is whether an agent removes friction that matters to the buyer. In the second, we went inside the enablement tier, where the near-term capital is moving, and laid out the four problems (access, identity, settlement, governance) that must be solved before anything above them works at scale.

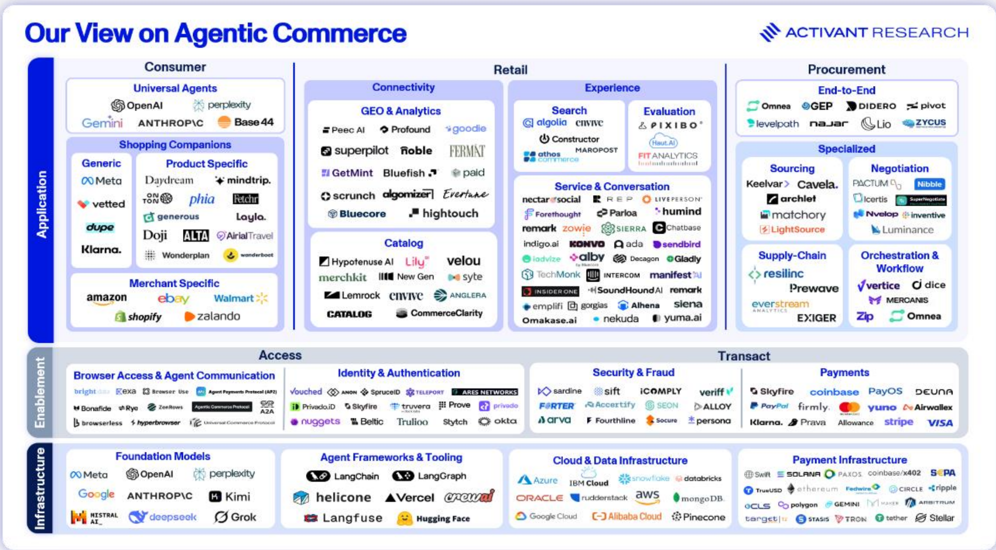

This piece moves up one layer. When the plumbing is built, and consumers start to delegate real purchases at real volume, how do brands and the companies serving them compete? We believe that the playbook is not uniform. It depends on what you are selling, and for some brands and merchants it is almost the opposite of the playbook that worked for the last 25 years.

One observation is worth making at the start. Most of the writing on agentic commerce that is currently circulating focuses on two things: what the big universal agents (ChatGPT, Gemini, Perplexity) will do, and how big the market will be. Both are interesting. Neither is very useful to a brand operator on a Tuesday morning trying to decide whether to invest in Generative Engine Optimization or ship a new product category. What is useful is a specific read on what changes at each stage of the buying funnel, and what a brand's quadrant says about which agents and playbooks will matter most. That is what we have tried to build here.

What Changes at Every Stage

Three shifts are happening across the buying funnel regardless of which category a brand sits in. They are what the agent layer does to the mechanics of commerce itself.

The first shift is **discovery**. For 25 years, brands invested in making sure their pages appeared at the top of search results. Agents do not work that way. An agent does not scan a results page and click a link. It queries structured product data directly, compares attributes, and decides without a human ever seeing the intermediate steps. When Adobe reported that AI-driven traffic to retail websites jumped twelvefold between July 2024 and February 2025, the headline was the growth.[1] The more important observation was the composition of that traffic. It behaves differently. Salesforce measured that traffic referred by AI agents converts at roughly eight times the rate of traffic from social media platforms, because the evaluation has already happened upstream.[2] The agent did the work that social platforms leave to the consumer. What the brand sees at the doorstep is pre-qualified intent, not curiosity.

The implication for brands is that the optimization target shifts from page rankings to something closer to shelf placement at the API layer. Generative Engine Optimization (GEO) is an emerging discipline, and it is not a rebrand of SEO. It is structurally different work. Product data needs to be structured cleanly, with consistent attribute formatting across every channel. Missing attributes are not minor technical gaps. They are reasons an agent will default to a competitor. Agents also draw on what has been written about a brand across the web. Third-party mentions, reviews, forum discussions, and editorial coverage all feed into how AI systems form recommendations. According to Noble, 95% of sources for mid-funnel searches are third-party websites, which reframes the brand visibility problem entirely.[3] The brands that will do best in the agent era are the ones that treat their catalog data and third-party footprint as core product, not as marketing overhead.

A new tooling category is forming around this. GEO and analytics players like Peec AI and Profound help brands understand how they are represented inside AI-generated responses and where their product data falls short. Noble takes a different approach, automating outreach and negotiation for brands to get included in sources that are used to train AI models. For consumers arriving onto merchant websites, search tools like Algolia and Constructor have moved toward intent-based search that understands context, not just terms. On-site search should be treated as a product, not a utility.

The second shift is **evaluation**. The hardest part of any purchase is the step before the click: deciding whether this product, from this brand, is right. For agent traffic, the evaluation problem is catalog completeness. An agent's context is based on the criteria it can access. If a brand's product data does not cleanly surface those attributes, the agent defaults to a competitor whose data does. For direct traffic, the evaluation challenge is similar but different. Consumers arrive with intent but often leave without converting because they cannot evaluate product fit. Will this fit, look right, suit my skin? That uncertainty is especially acute in fashion, footwear, and beauty, where specs alone do not settle the question. Category-specific evaluation tools are emerging to close the gap. Haut AI focuses on personalized skincare analysis. FitAnalytics handles size and fit guidance. Pixibo offers a similar tool, claiming to increase conversion rates by 20% to 30% while decreasing return rates by 3% to 4%.[4] If those numbers are even directionally accurate, the operating leverage is significant, and category-specific evaluation tools become one of the more structurally interesting in the application tier.

The third shift is **conversion**. For decades, conversion optimization meant reducing friction at checkout. Agents reduce that friction to nearly zero. The new competitive edge is the last step before conversion, where a customer on the fence either commits or drops out. Service agents, deployed at the point of conversion rather than post-purchase, are the emerging answer. 44% of online consumers consider having their questions answered during a purchase to be one of the most important features a website can offer. Visitors who engage with live chat are 2.8 times more likely to make a purchase.[5] Tools like Sierra, Forethought, and Gorgias operate as AI-powered service layers that clarify shipping, offer contextual discounts, and handle return queries in real time. Gorgias claims its AI shopping assistant can drive up to 50% higher engagement and 20% more revenue per visitor.[6] Salesforce's data found that retailers with branded AI agents experienced seven times the sales growth compared to those without.[7] Service is no longer a post-purchase cost center. It is a pre-purchase conversion tool.

Those three shifts apply to every brand in every quadrant. What differs by quadrant is what a brand should prioritize, which agent category will dominate its part of the market, and what the unit economics of competing look like. The four playbooks below walk through each quadrant in turn.

The Commoditized Quadrant Playbook

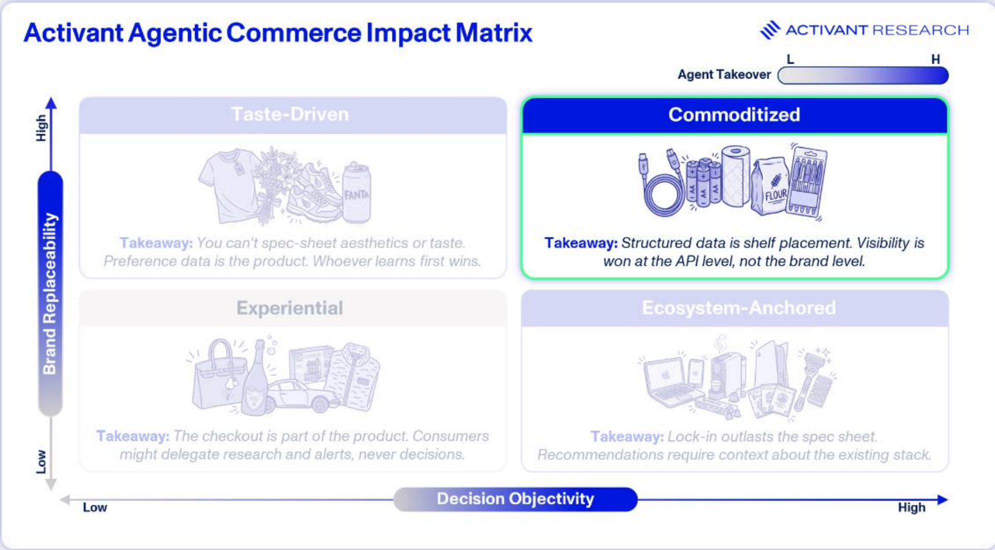

In the commoditized quadrant, the structural logic favors scale, and the playbook for brands and agents is the most settled of the four. The largest universal agents (ChatGPT, Claude, Gemini, Perplexity) and the merchant platforms with inherent data advantages (Amazon, Walmart) have a head start that is difficult to close. Walmart attributing 20% of its referral traffic to ChatGPT is what distribution looks like when the major retailers commit.[8] If your product lives in this quadrant (USB cables, batteries, printer paper, commodity household goods), the near-term economics favor being visible on those surfaces, not building an independent direct channel.

For brands, the playbook has three moves. First, invest in GEO and catalog enrichment as core operational infrastructure, not as a marketing experiment. Second, accept that branding-as-differentiation works less well in this quadrant than it used to, and compete on the structured attributes the agent can see (price, delivery speed, reliability, review score). Third, use subscription and auto-reorder as the moat, because a consumer who has already delegated a category is unlikely to take it back. The cost of winning the first agent-mediated purchase is high. The cost of keeping the second, third, and fourth is much lower.

For agent-builders, the near-term returns in this quadrant accrue to the GenAI pioneers and the merchant platforms. The bar for a new entrant competing head-to-head is unusually high. The question worth asking is whether your wedge is a feature the pioneers will eventually ship, or a category they structurally cannot build. In the commoditized quadrant, most of what we see falls into the first bucket.

The exception is catalog enrichment. The recommendation quality of every major platform like ChatGPT and Gemini depends entirely on the quality of the structured data merchants publish. Most merchant catalogs are messy: inconsistent attributes, missing specs, poor taxonomy, thin descriptions. An agent cannot recommend a product it cannot parse. Catalog enrichment platforms like Lily AI, Catalog, CommerceClarity, and Hypotenuse AI are building the infrastructure that determines whose products get surfaced and whose get skipped. They do not compete with the scale layer. They are the reason the scale layer works. That is a structurally defensible position, and it is where we look first in this quadrant.

> *The retailer describes it as midnight french terry athleisure. The customer searches for a navy hoodie. That gap is not a minor inconvenience - it's lost revenue at scale, and most merchants can't even measure it. Multiply that by every product, every query, every surface where an agent is making a recommendation, and you start to see why catalog data isn't a back-office problem anymore. It's the determining factor in whether you show up at all."*
>
> Purva Gupta
> Co-founder and CEO, Lily AI

The Ecosystem-Anchored Quadrant Playbook

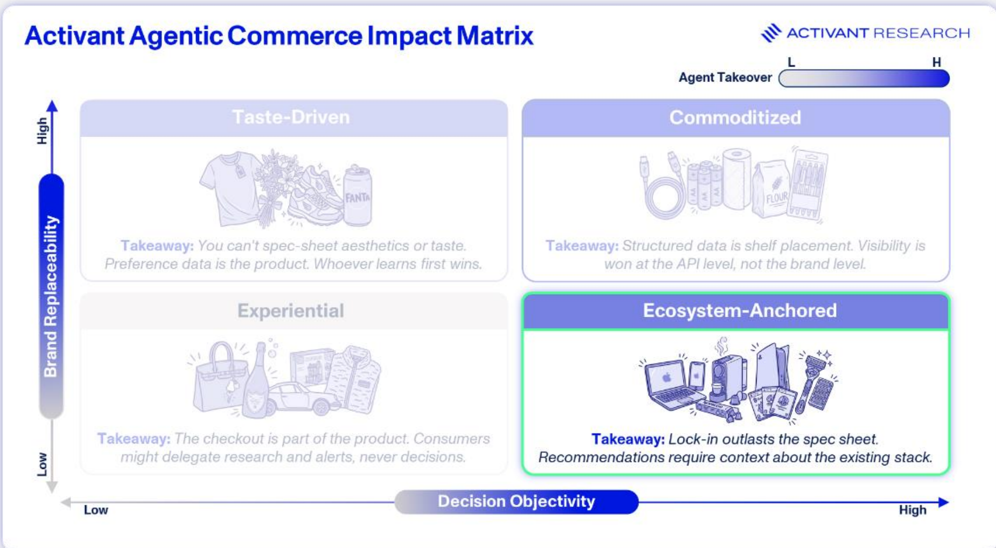

In the ecosystem-anchored quadrant (laptops, smartphones, smart home devices), the space is largely unresolved. The specs are measurable, which should favor agent mediation, but brand replaceability is low because of switching costs. An iPhone user may find a Windows laptop that scores higher on every benchmark at a lower price, but iCloud and AirDrop make switching a poor choice. Any agent operating in this quadrant needs to understand the user's existing stack, not just compare products in isolation.

Very few companies are building agents specifically for this quadrant today. We think that is a gap, not a reflection of market size. The first independent agent that reliably understands ecosystem dependencies will have a real product. It is, however, harder to build than it sounds. It requires context that product catalogs do not contain, which is precisely why the incumbents with the deepest user context are better positioned than new entrants to solve it.

Two kinds of players are well positioned to claim it. The first is the major GenAI platforms, if they extend their memory models far enough. Not just broad user preferences, but granular stack awareness. A user who asks how to adjust settings on their iOS device is clearly in the Apple ecosystem and should be steered toward a Mac over a Windows PC when they next need a laptop. A user with a Sonos system who asks which speaker to add to their bedroom should not receive a recommendation for a standalone Bluetooth speaker that scored higher in a review, even if the specs support it. The agent that knows the stack gives the useful answer; the agent that ignores it gives the technically correct but practically wrong one.

The second player type is personal data platforms with access to rich contextual signals: Meta via WhatsApp and Messenger, Apple via iMessage. Understanding a user's ecosystem and life circumstances requires understanding their personal context, much of which is reflected in conversations with friends and family. Whether users would allow that level of access, and whether regulators would permit it, is genuinely open. But the data advantage available to whoever solves that problem is substantial. Realizing that advantage will also require agents to clearly communicate to users why, given their personal circumstances, a certain product scores better than others.

For brands in this quadrant, the playbook is defensive. Deepen the ecosystem. Make switching costs more real, not less. Make sure that agents understand your ecosystem benefits, by being part of the community discussions explaining them. These discussions, whether in Reddit threads, YouTube comments or elsewhere, part of the training data that shapes an agent's recommendations. For agent-builders, the playbook is the opposite. Find the narrow space where cross-ecosystem evaluation matters and the incumbents cannot credibly build it, and build there first. It is a smaller near-term opportunity than the taste-driven quadrant, but it is more defensible because the structural reason that the pioneers cannot compete is unlikely to change.

The Taste-Driven Quadrant Playbook

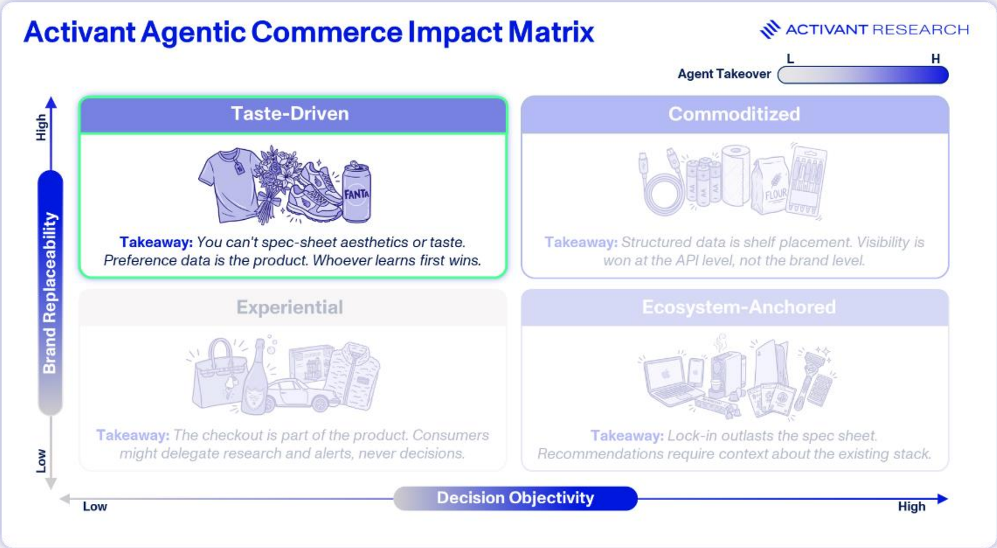

The taste-driven quadrant is where the most interesting and most contested company-building is happening in agentic commerce, and it is also where the playbook is least settled. Consumers are open to switching brands, but the decision runs on personal aesthetics that no spec sheet captures. Instagram, TikTok, and Pinterest hold years of visual-taste signal, captured in what users pause on, save, and return to. The commercial proof of how much that signal is worth is already visible: TikTok Shop reached an estimated $33B in GMV in 2024, driven largely by discovery-led purchases in fashion, beauty, and home goods. The platforms that already hold that data appear to have an obvious head start.

What we have found interesting is how slowly that advantage has converted. TikTok Shop has scale, but it has not closed the loop between taste signal and confident purchase recommendation in the way its data would suggest it should.[9] The gap between "we know what this user pauses on" and "we can recommend a specific product that fits their taste well enough to drive a purchase without seeing returns cannibalize the economics" turns out to be wider than the underlying data implies. Passive signals tell you what a user has engaged with. They do not tell you what they will be happy to buy, especially in categories where size, fit, and physical presence matter.

That gap has created space for independent companies building around different signal architectures, and it is the part of the application tier we find most analytically interesting right now. These companies are comparing products based on appearance and preference instead of spec sheets and price alone.

Daydream and Onton are two early examples, both working with recommendation quality as the primary wedge. Other approaches we are watching use explicit preference elicitation rather than passive signal, treating the consumer as a co-author of the recommendation rather than a data source. Others are building around community curation, where the signal is not what the user themselves engages with but what the users they trust engage with. Others are starting from specific subcategories (athletic footwear, beauty, home furnishings) where the platforms have paid less attention and where narrow-vertical data compounds faster than broad-vertical data.

We do not yet know which architecture wins. We think the question is more open than consensus suggests. The companies we find most compelling in this quadrant are not trying to out-scale the platforms. They are trying to out-understand them in a narrower vertical, where the depth of signal matters more than the breadth. The brands that will do best in this quadrant are the ones that recognize the difference between platforms that hold taste data and platforms that drive taste-based purchase, and who invest accordingly. Building an Instagram presence is not the same as building a presence on whichever taste-aware agent wins the next five years, and the brands treating them as interchangeable are going to learn an expensive lesson.

For brands, the playbook has three moves. First, invest in first-party data collection around aesthetic preference, fit, and taste, because the taste-aware agents are going to need that data and the brands that already have it will gain leverage. Second, treat evaluation tools (fit, personalization, visualization) as core conversion infrastructure, not as post-purchase support. Third, build real partnerships with the independent taste-aware agents early, because the ones that work are going to need co-development on product data, preference models, and return handling, and first-mover partnerships here will matter. For agent-builders, the returns in this quadrant are not yet decided, and the firms that deploy capital early against architectural differentiation will have a structural advantage if the independents do emerge as the winners.

The Experiential Quadrant Playbook

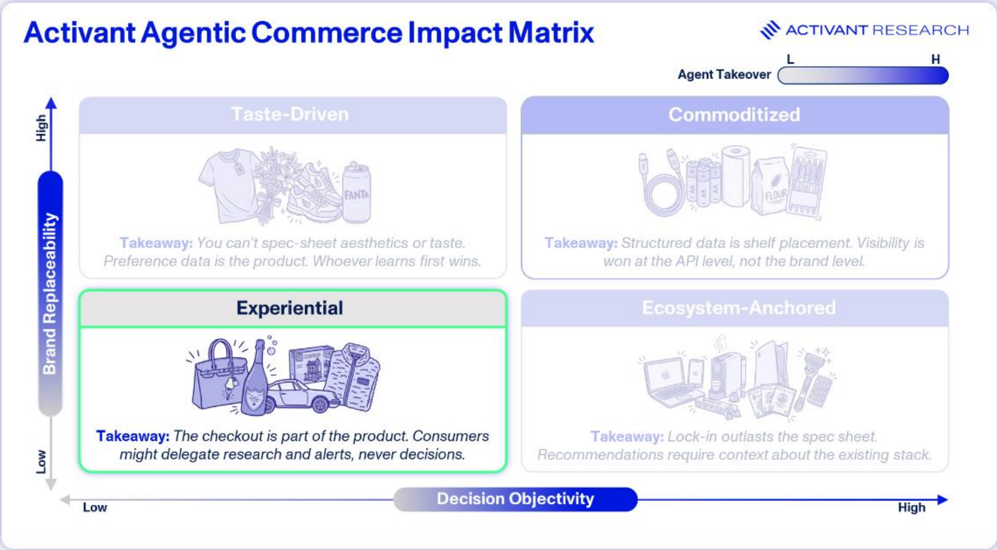

In the experiential quadrant (luxury fashion, fine jewelry, premium watches), the structural logic works against agent mediation. Both axes work against it. The brand carries meaning that no comparison can capture. An agent cannot evaluate a $15,000 Rolex and recommend a Casio that looks similar, because the purchase was never about timekeeping. Richemont's 6% online penetration, despite 25 years of e-commerce maturity, is the cleanest evidence we have that this quadrant moves slowly and for reasons agents cannot remove.[10]

For brands in this quadrant, the playbook is the simplest of the four: do not compete on agent-friendly terms. Brands that try to dilute the very distinctiveness that protects them to win agent-mediated transactions trade long-term moat for short-term volume, and the trade rarely pays. Agents play a supporting role in this quadrant at best. The companies creating value here are using agents for concierge-style assistance, authentication, inventory discovery, and post-purchase service, rather than as the transaction layer. Sierra, Forethought, and Gorgias are live examples in the service layer, and the economics we flagged above (Gorgias's 50% engagement lift, Salesforce's 7x sales growth) are available to experiential brands on the direct-traffic side regardless of whether they participate in agent-mediated transactions.

The strategic question for an experiential brand is not whether to prepare for agents. It is how to use agents to deepen the direct relationship with the consumers they already have, rather than to compete for new transactions in a channel that structurally does not value what makes their brand valuable.

Market to People, Sell to Software

The retail industry has now been through this pattern twice. First with e-commerce. Now with agentic commerce. The lesson from the first transition was that the honest outcome was not replacement. It was coexistence in most categories, near-total displacement in a few, and almost no change in a handful. The lesson is not that software beats physical retail. It is that different categories moved at different speeds, and the brands that survived were the ones that read their category correctly and moved accordingly.

Agentic commerce is the same transition at a different layer of the stack. Some categories will move to agents almost entirely. Some will shift partially. Some will barely move at all. The brands that will do best are the ones that recognize which quadrant they occupy and play to its logic, rather than applying a single playbook to all four.

For the last 25 years, e-commerce operators optimized for human attention: ad placements, landing pages, search rankings, cart abandonment flows. Those levers still work, and they will continue to work for the traffic that arrives directly at brand sites. But a growing share of transactions will originate on surfaces where the buyer is no longer human. Catalog data structured for agent consumption, third-party coverage that shapes how AI models represent a brand, verifiable product attributes that surface in agent comparisons, and service layers that answer agent-mediated queries fast enough to preserve conversion: these are the new levers. They are not improvements on the old marketing stack. They are a parallel stack.

Brands need to market to people. Increasingly, they also need to sell to software. That is not a clever framing. It is the operational fact that the next decade of retail will run on. The operators and investors who do well will be the ones who recognize that both channels are real, that each has its own mechanics, and that treating them as the same channel with different surfaces is the fastest way to lose the transition.

* * *

The question for the next decade of retail is not whether agents change commerce. They already have. The question is who reads the shape of that change correctly, and who keeps applying the wrong playbook to the wrong quadrant until the market corrects them. If you are building or investing in this space, we would love to hear your perspective. Where do you believe agents will create the most value? Feel free to share any comments or thoughts by reaching out.

---

[1] IMD, Generative Engine Optimization (GEO), November 2025

[2] CMS Wire, Cyber Week 2025, December 2025

[3] Noble, Company Website, Accessed April 13 2026

[4] Pixobo, Company Website, Accessed April 8 2026

[5] ibid

[6] Gorgias, Company Website, Accessed April 9 2026

[7] Salesforce, 2025 Cyber Week Predictions, November 20 2025

[8] Digiday, ChatGPT is now 20% of Walmart's referral traffic, September 25 2025

[9] ECDB, Gross Merchandise Value of TikTok Shops worldwide in 2024, January 2026

[10] HSBC Global Investment Research, Global Luxury Goods, March 2026