

# Agentic Commerce

A New Channel  
Not a New World

June 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.

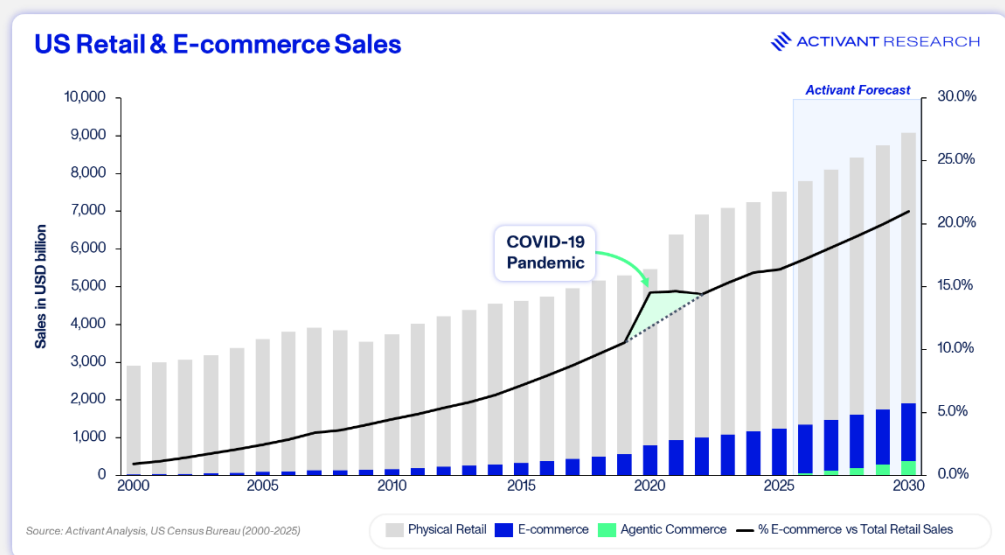
You can find out more about Activant and our research at <https://activantcapital.com/>.

The conversation around agentic commerce has a blind spot. It assumes a single, inevitable wave. Agents will take over shopping. Brands will lose relevance. Consumers will hand over their wallets to software. It makes for good headlines, but we believe it gets the story wrong.

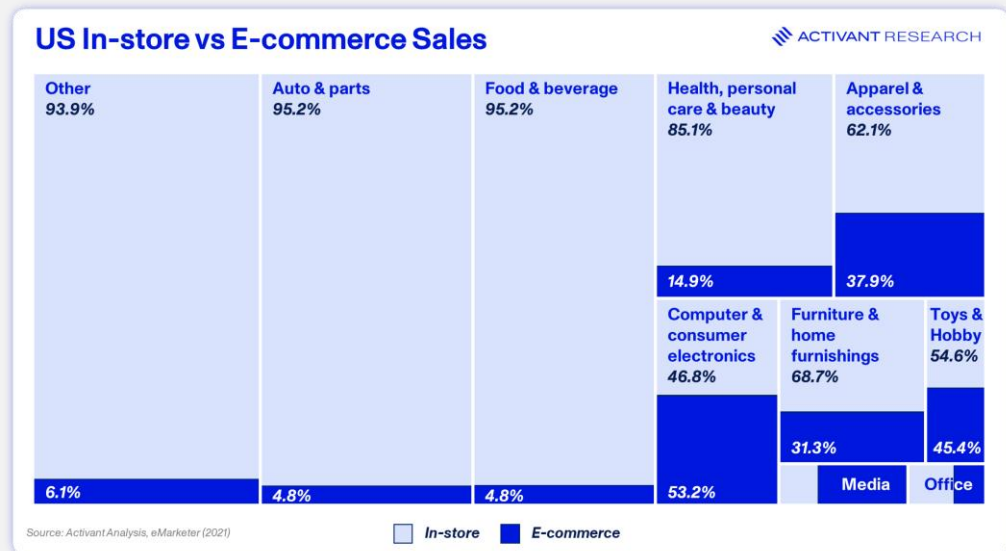
The reality is messier and more interesting. We believe agents will dominate certain categories of commerce but will be less relevant in others. The line between those two outcomes has nothing to do with technology itself, but rather whether agents can remove friction that matters to the buyer.

Most people don't care which company made their USB cable. Many care about which watch is on their wrist. An agent that picks the cheapest, highest-rated AA batteries with next-day delivery might do you a favor. An agent that independently orders your next car is likely overstepping. This distinction sounds obvious when you say it aloud, but almost nobody in the agentic commerce conversation is building frameworks around it. Instead, we get sweeping market projections with wide ranges. Morgan Stanley, for example, suggests that agentic shoppers could represent \$190 to \$385 billion in U.S. e-commerce spending by 2030<sup>1</sup>, while McKinsey estimates that global figures might run as high as \$3 to \$5 trillion.<sup>2</sup> These numbers may be directionally right, but we think they obscure the question that matters: where does agent-mediated commerce create enough value to change buying behavior and who is building for this future?

The right historical comparison is e-commerce itself. Online retail did not replace physical stores. Twenty-five years in, US e-commerce still only accounts for roughly one in six retail dollars, and the annual growth curve is a steady gradient rather than a cliff.<sup>3</sup>



That is not failure. It is what category-by-category penetration looks like at scale. Most media sales, including books, music, and video moved online, reaching 69.1% online sales penetration in 2021.<sup>4</sup> Apparel and accessories, the category with the fourth-largest penetration, sits at 37.9%, with brand and category composition driving the spread.<sup>5</sup> Richemont, the luxury group that owns Cartier, IWC, and Montblanc, for example, still only generates 6% of its sales online.<sup>6</sup> Categories like auto and parts or food and beverage barely moved at all, achieving only 4.8% e-commerce sales.<sup>7</sup>



Agentic commerce will likely follow the same path: fast penetration where friction is high and decisions are objective, much slower penetration where the buying experience carries its own value. The two channels will coexist.

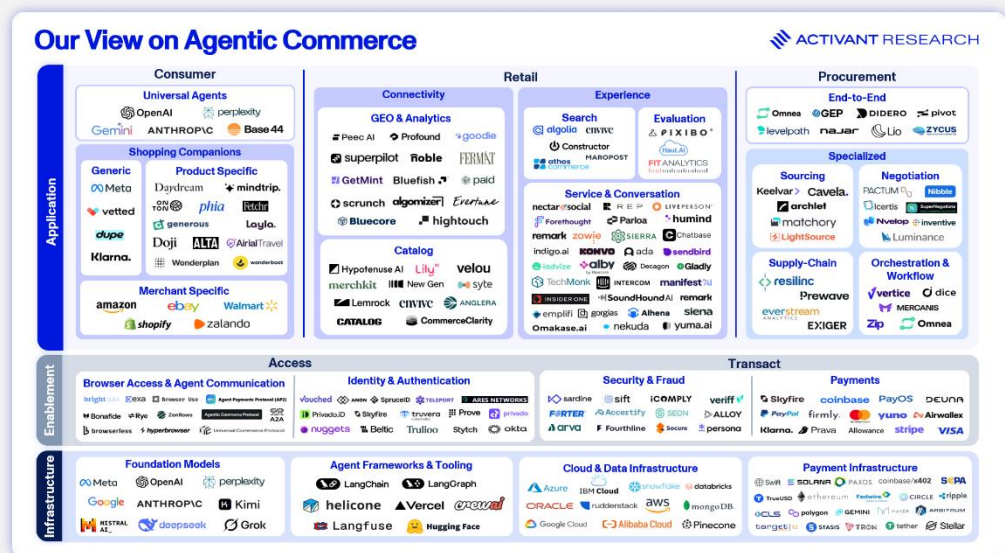
Over the past several months we have studied how agentic commerce is developing across the consumer landscape. We mapped adoption patterns by product category and tracked how different stakeholders are positioning. We also stress-tested the infrastructure being built underneath. The deeper work was understanding the spectrum of agentic commerce and what is required for agents to discover, recommend, and transact.

Right now, agents mainly recommend but are seldom able to transact autonomously. Even the largest players have struggled with this. [OpenAI](#) launched in-chat checkout in September 2025, then shut it down within six months, citing a lack of flexibility. It has since pivoted back to product discovery, redirecting users to retailer's sites to complete purchases themselves.<sup>8</sup> [Perplexity](#)'s Instant Buy feature still works only for merchants integrated with [PayPal](#)'s checkout infrastructure, and for US customers only.<sup>9</sup>

This is opening space for other entrants. [Allowance](#), for example, is building payments infrastructure and has also launched a consumer app. The app lets any customer set up agents to buy from any website under set constraints like price, merchant, and time. It hides credit card details from the agents while still collecting rewards, and it keeps an audit log of every purchase.

## The Stack, and Where Value Accrues

Agentic commerce is not one market. It is a stack of three layers with multiple subsectors. Understanding the solutions available in each helps explain why certain product categories will naturally benefit more than others.



At the top sits the **application layer** where consumers, retailers, and procurement teams interact with agents. Universal agents like [ChatGPT](#), [Claude](#), [Gemini](#), and [Perplexity](#) offer a single platform to users already using LLMs for other purposes. Merchant-specific agents from [Amazon](#), [Walmart](#), and [Shopify](#) have a distinct advantage: they already hold the structured product, pricing, and inventory data that agents need, and can integrate agentic capabilities directly into existing purchase flows. Shopping companions like [Daydream](#), [Onton](#), and [Phia](#) are specializing in product categories where universal agents struggle because they compare products on spec sheets rather than visual appearance. [Layla](#) focuses on travel, a category that requires an understanding of experiences, accommodations and flight connections simultaneously.

Retailers are also building out their own agentic toolkits, and a distinct vendor landscape is forming around two problems: how to reach agents operating outside the merchant's site, and how to improve the AI experience on it. The first category covers connectivity. Generative Engine Optimization (GEO) and analytics tools like [Peec AI](#), [Profound](#), and [Algomizer](#) tell retailers how often and how accurately AI systems surface their products when shoppers ask an external agent for recommendations. Cataloging tools like [Catalog](#) and [Hypotenuse AI](#) take a complementary approach: they structure and enrich product data so agents can parse, compare, and rank it more reliably. One measures discoverability; the other improves it.


The second category covers on-site experience. Search tools like [Algolia](#) and [Constructor](#) replace keyword matching with conversational, intent-based queries. Fit and evaluation tools like [FitAnalytics](#) and [Haut AI](#) help customers resolve practical questions about size, type, or compatibility before buying. And service tools like [Sierra](#) and [Forethought](#) handle customer conversations through agentic support flows. Together, these aim to make a retailer's own website feel as responsive as the best external agents.

Corporate procurement is a parallel story. A growing set of agentic tools now helps businesses buy at scale, but that market has different dynamics and different players. This is why we are setting it aside for this piece to focus on the retail market.

The middle layer of the agentic commerce stack is **enablement**. Access, identity and authentication, security and fraud, and payments tools sit here. This is the plumbing that determines whether any application-layer agent can find relevant products or complete a purchase. Browser access tools like [Browser Use](#) and [Browserless](#) let agents navigate merchant sites without any merchant cooperation, at the cost of reliability. The API-native alternative is being built around two competing protocols: the [Agentic Commerce Protocol \(ACP\)](#) from OpenAI and Stripe and the [Universal Commerce Protocol \(UCP\)](#) from Google, co-developed with Shopify, Walmart, and Target.

Identity and authentication tools like [Okta](#) and [Trulioo](#) provide Know Your Agent (KYA) solutions, enabling agents to prove that they are acting on behalf of a verified individual. At the same time, security and fraud infrastructure is also being rebuilt. Current systems rely heavily on human behavioral signals to flag bad actors: typing patterns, mouse movement, session timing. Agent-initiated transactions strip all of that out, making them look nearly identical to fraud. [Sardine](#) (an Activant portfolio company) and [Forter](#) (an exited Activant portfolio company), alongside [Sift](#), are adapting their architectures to tell the difference between a bot that is stealing and a bot that is buying.

Payment rails need to adapt too. [Visa](#) and [Mastercard](#) have both unveiled agent credentialing frameworks built on Web Bot Auth, a Cloudflare-proposed protocol undergoing standardization at the [IETF](#) that lets merchants verify an agent's identity before a transaction clears. Allowance is building a platform that provides agents with guardrails for spending, limiting amounts and merchants with time bound payments without exposing card details. [Stripe](#), [PayPal](#), and [Coinbase](#) are moving in the same direction. With several agent-friendly solutions emerging but no clear winner yet, intelligent payment infrastructure is gaining importance. [Deuna](#), another Activant portfolio company, built a solution that allows merchants and financial institutions to not only optimize their commerce KPI's, but to turn their previously unused or convoluted payments data into strategic insights through their Athia platform. All by utilizing a Payments Agentic Workforce consisting of multiple specialist agents focused on maximizing a business's overall P&L.



*Deuna solves the complex, undiscovered problems that constrain enterprise ROI and erode margin. Powered by Athia and the client's data, our Payments Agentic Workforce delivers everything from strategic insights and direction to real-time decisions, all individually tuned to their business."*

**Roberto Kafati**  
Co-Founder and CEO, Deuna

The bottom layer is **infrastructure**, the furthest upstream and the most mature. Foundation model providers OpenAI, [Anthropic](#), [Google](#), and [Meta](#) compete at the top, with open-source alternatives from [Mistral](#), [Deepseek](#), and [xAI](#) pushing the frontier. Agent frameworks like [LangChain](#), [Vercel](#), [CrewAI](#), and [Hugging Face](#) sit one layer below. Cloud and data infrastructure is already an oligopoly of [AWS](#), [Azure](#), [Google Cloud](#), and [Oracle](#), with [Snowflake](#), [Databricks](#), and [Pinecone](#) serving more specific data workloads. Traditional payment rails like Swift, SEPA, and Fedwire are being supplemented by agent-native settlement layers from [Circle](#), [Paxos](#), Coinbase's [x402](#), and the stablecoin protocols built on [Ethereum](#), [Solana](#), and [Tron](#). Our view is that this is not where the next generation of agentic commerce returns will be made. The layer is too contested, the capital requirements are too high, and the companies that will dominate it are largely already public or on their way. Returns in this cycle will likely accrue to what gets built on top.

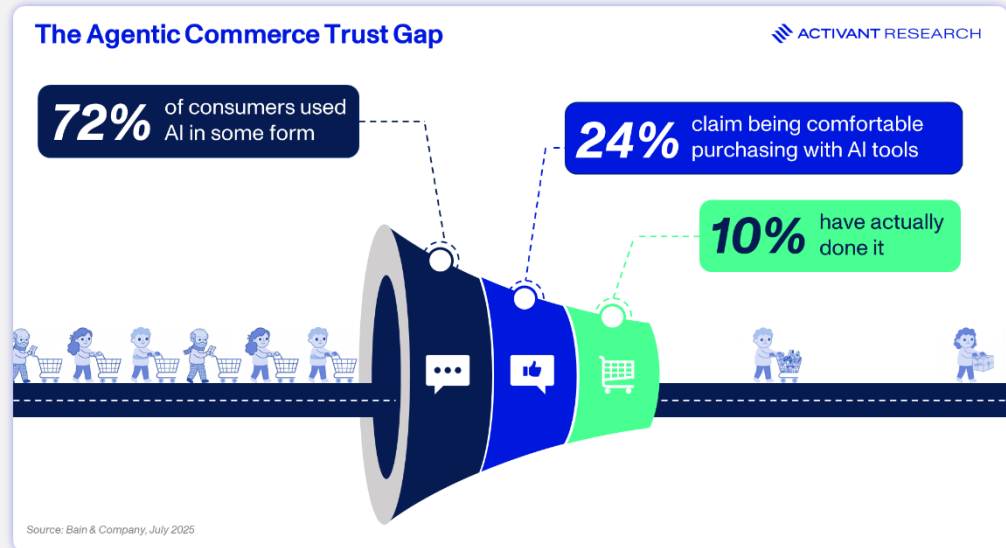
That leaves application and enablement as the two layers worth close attention. They behave differently. Application is where the eventual franchise-scale businesses get built, but most depend on enablement functioning effectively. Enablement is where capital is flowing because every application-layer agent needs it, and large parts remain unsolved. In equity terms, enablement is a near-term bet on a layer that must exist. Application is a medium-term bet on which agents consumers will adopt, and this will likely evolve over the next 18 to 36 months.

The stack does not move on a single clock. Enablement is being built right now. Application is being experimented with. Governance is lagging everything. The Amazon lawsuit against Perplexity in late 2025 showed what that lag looks like in practice.<sup>10</sup> A browser agent disguised itself as a human-operated Chrome session to complete purchases inside Amazon's checkout flow, bypassing the controls Amazon had built to detect automated behavior. It was less a sophisticated attack than an opportunistic workaround, and that is what made it revealing. The authorization question, which had seemed abstract in 2024, became concrete and litigious almost overnight. Until courts settle whether user consent legally extends to agents acting on their behalf, and who is liable when an agent purchase goes wrong, the infrastructure will be ready before the rules are. That gap is where adoption lags what is technically possible. Reading those timing differences correctly is where the analytical work sits, more than picking winners.


## How Agents Create Real Value

Consumer behavior data tells a clearer story than the forecasts do. In July 2025, Bain & Company published a survey on US consumers' use of AI-assisted shopping. Seventy-two percent said they had used AI in some form. Twenty-four percent said they were comfortable purchasing with AI tools. Only ten percent had done it.<sup>11</sup>

That gap between 72 and 10 is the entire piece in one number. Most consumers know about agents. A minority say they would use one to buy. Fewer still have. The constraint is not capability. It is trust, and trust moves on its own timetable.



What the research shows clearly is how trust extends. 89% of consumers who use AI for shopping still double-check the recommendation before buying, and 88% said knowing where the recommended information came from would make them more confident.<sup>12</sup> Consumers are not asking for perfection. They are asking for transparency. Agents that show their work earn delegation faster than those that skip straight to the answer.



*Most online travel agencies used to just forward you somewhere; you didn't even need to log in. Now the bar is completely different. The biggest travel company of the next decade will not be the one with the most inventory. It will be the one that knows you best."*

**Saad Saeed**

Founder and CEO, Layla AI  
Previous Co-Founder, Flink

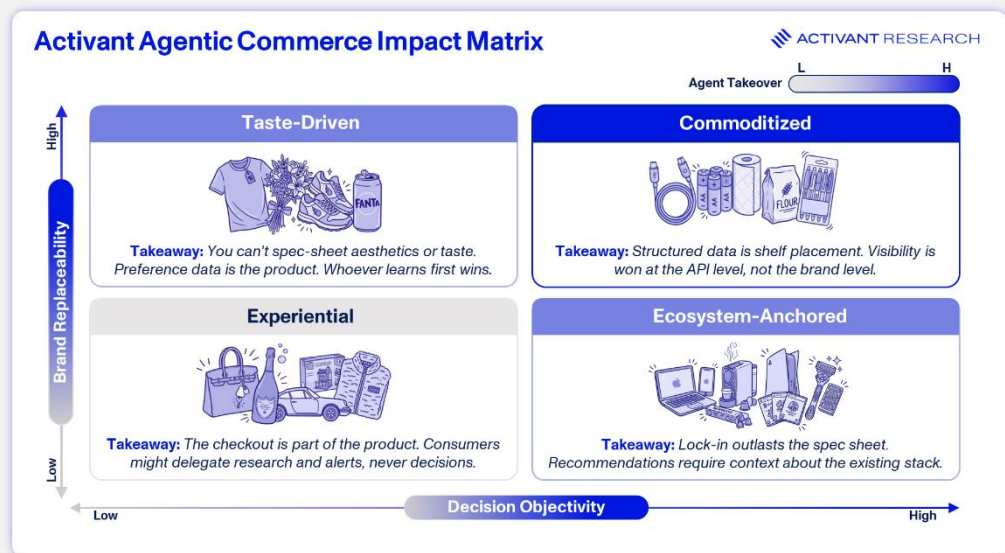
Brand also matters when it comes to trust. Bain's research shows that consumers are more than twice as comfortable delegating to high-trust consumer brands, such as Apple, PayPal, or Google than banks or credit card companies.<sup>13</sup> Familiarity compounds, and in this category, it matters more than product specifications.

The good thing about trust is that it can be built step-by-step by creating clear customer benefit. In 2024, Americans redeemed roughly 871 million coupons.

That is only 1.3% of coupons issued.<sup>14</sup> Consumers know the savings exist. Many do not collect them because the effort cost is too high. Agents could help to remove that cost. A consumer who would never open a separate tab to check a cashback site gets the benefit automatically because the agent handles it as part of the purchase flow. This is not a marginal improvement in existing behavior. It is latent demand turning into captured demand the moment the friction drops to zero. It is also a major threat for retailers and consumer goods companies used to dealing with very low coupon redemption rates.

Where the trust gap is already closing, the benefits are significant. Adobe reported AI-driven traffic to retail websites jumped twelvefold between July 2024 and February 2025.<sup>15</sup> Walmart attributes 20% of its referral traffic to ChatGPT.<sup>16</sup> Salesforce measured that traffic referred by AI agents converts at roughly eight times the rate of traffic from social media platforms.<sup>17</sup> Where consumers are ready to act, the commercial signal is real. For tool builders and merchants alike, the question is which categories see that readiness next.

To answer that question, we mapped the consumer products space along two dimensions. The horizontal axis captures decision objectivity. Can the purchase be made using structured specifications like price, storage capacity, wattage, and delivery speed, or does it depend on subjective factors like design, taste, and emotional resonance? The vertical axis captures brand replaceability. How easily can a consumer substitute one brand for another, and how willing are they to do so? Four categories emerge.



**Commoditized** products (USB cables, batteries, printer paper) sit in the top right. An agent compares them on price, ratings, and delivery speed in seconds, and most consumers do not care who made them. The Adobe and Salesforce numbers cited above are mostly this quadrant. If your product lives here, you are only as visible as your structured data. GEO is the new shelf placement, and the companies building around it, including [Peec AI](#), [Profound](#), and [Noble](#), are early but not speculative. Noble reports that 95% of sources for mid-funnel searches are third-party websites, which reframes the brand visibility problem entirely.<sup>18</sup>

The structural logic here favors scale. The largest universal agents and platforms with inherent data advantages, ChatGPT, Claude, Amazon, and Walmart, have a head start that is difficult to close. Walmart attributing one-fifth of its referral traffic to ChatGPT is a preview of what distribution looks like when the giants commit. That does not mean the quadrant is settled, but the bar for new entrants is high. The companies worth watching are solving specific sub-problems: catalog enrichment players like [Catalog](#), [CommerceClarity](#), and [Hypotenuse AI](#), and on-site search players like [Algolia](#) and [Constructor](#), rather than competing head-to-head with the platforms.

**Ecosystem-anchored** products (laptops, smartphones, smart home devices) sit in the bottom right. The specs are measurable, 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 the worst choice. Agents working here need to map the user's existing stack, not just compare products in isolation. The defenders are structurally strong. The more interesting question is whether new agents can surface cross-ecosystem decisions, comparing a Windows laptop against a MacBook in a way that accounts for both performance and lock-in, rather than trying to displace incumbents outright. Very few companies are building for this specific problem today. We think that is a gap, not a reflection of market size.

**Taste-driven** products (fast fashion, home decor, food and beverage) sit in the top left. Consumers are open to switching brands, but the decision runs on personal aesthetics that no spec sheet can capture. This is the most technically demanding quadrant for agents. [Instagram](#), [TikTok](#), and [Pinterest](#) hold years of visual-taste signal, and the commercial proof is visible: TikTok Shop reached an estimated \$33B in GMV in 2024, driven largely by discovery-led purchases in fashion, beauty, and home goods.<sup>19</sup> The platforms that hold that data appear to have an obvious head start. What we find 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.

That gap has created space for independent companies to build different architectures. [Daydream](#) and [Onton](#) are both working in this space with recommendation quality as the primary wedge. Others are using explicit preference elicitation, community curation, or narrow subcategory focus. Category-specific evaluation tools like Haut AI in skincare and FitAnalytics and [Pixibo](#) in sizing are doing related work on the evaluation side. Pixibo claims to increase conversion rates by 20 to 30% while reducing returns by 3 to 4%, which if accurate is the kind of operating-metric improvement that moves a full P&L.<sup>20</sup> We do not yet know which architecture wins. We think the question is more open than consensus suggests, and the companies we find most compelling are not trying to out-scale the platforms. They are trying to out-understand them in a narrower vertical.

**Experiential** products (luxury fashion, fine jewelry, premium watches) sit in the bottom left. Both axes work against agent mediation. Richemont's 6% online penetration is a clear data point: a category that had 25 years to go digital and largely did not.<sup>21</sup> 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.

Agents play a supporting role here at best, and that is where the investment logic sits: concierge-style assistance, authentication, inventory discovery, and post-purchase service through tools like [Sierra](#), [Forethought](#), and [Gorgias](#). Gorgias reports its AI shopping assistants can increase engagement by 50% and achieve 20% more revenue per visitor.<sup>22</sup> Salesforce found retailers with branded AI agents experienced seven times the sales growth compared to those without.<sup>23</sup>

The post-purchase service layer is a real investment category regardless of which quadrant the product sits in. Similarly, the enablement layer sits underneath all four quadrants and potentially represents the clearest near-term opportunity.

## Implications for Merchants, Investors and Founders

Every application-layer agent needs identity, fraud, and payment infrastructure to function before anything above matters. This is why the Activant portfolio is deployed here today with Sardine, Forter, and Deuna.

For merchants, the transition is not from mattering to not mattering. It is from managing one channel to competing across two simultaneously. The second

channel does not replace the first. It runs in parallel, on different logic, rewarding structured data and transparency over brand spend and search optimization. That is a harder game. It is also a more honest one.

For investors, the framework above suggests different return profiles by quadrant. Commoditized categories are a scale and data game where early movers in catalog enrichment and GEO monitoring have a real window, but it closes fast. Taste-driven categories are more architecturally open and the outcome less settled, which means more risk and more upside. Ecosystem-anchored categories are underbuilt relative to the potential size. Experiential categories are a service and support story, not a purchase transaction one.

One number frames the stakes. If agents mediate only one in five e-commerce transactions by 2030, that is still roughly \$300B to \$400B in volume, or about 4% of US retail sales. But that market only materializes if consumers hand over transactions, which brings it back to the trust gap. Only 10% of consumers have made an AI-assisted purchase. The curve from 10% to 50% will run through these quadrants at very different speeds, and reading which quadrant moves when is where the timing analysis matters most.

## A New Channel, Not a New World

The e-commerce penetration data presented earlier in this article tells the story of the first transition. Every category moved at its own pace, and the pace was set by how much of the purchase experience could survive translation to a screen.

Agentic commerce is that same transition running a second time, one layer deeper in the stack. The categories will not repeat perfectly. The shape of the outcome will. Some categories will move to agents almost entirely. Some will shift partially. Some will barely move at all. The operators and investors who do well here will likely be the ones who read their quadrant correctly and move at the pace each category allows. The categories that look agent-friendly on paper are not always the ones where agents create durable value. The spec sheet is a tempting guide. It is not a complete one.

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.

Three dynamics will determine the pace. First, trust infrastructure. The payment networks, identity layers, and agent credentialing frameworks discussed above are prerequisites, not features. Until they mature, agent-completed transactions will remain concentrated in low-risk, commoditized categories. Second, data readiness. The merchants that structure their catalogs, pricing feeds, and inventory APIs for machine readability will be visible to agents; the rest will not. Third, regulatory clarity. Agent liability, consumer consent, and cross-border data governance are open questions today. How they resolve will shape which business models survive and which get regulated into irrelevance.

\* \* \*

This first article mapped the application layer. The next in the series dives deeper into the enablement infrastructure that determines whether any of this will scale. The browser versus API question, identity and authentication, agent-native payments, and the governance and liability problem that sit underneath all three. That is where the structural bets are being placed right now, mostly out of sight. If you are building or investing in this space, we want to hear from you. Where in the stack do you think the value will accrue? Feel free to share any comments or thoughts by reaching out.

**Disclaimer:** The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at [activantcapital.com/companies/](http://activantcapital.com/companies/) for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.

---

<sup>1</sup> Morgan Stanley, [Here Come the Shopping Bots](#), December 2025

<sup>2</sup> McKinsey, [The agentic commerce opportunity](#), October 2025

<sup>3</sup> US Consensus Bureau, [Quarterly Retail E-Commerce Sales Q4 2025](#), March 2026

<sup>4</sup> eMarketer, [US Ecommerce by Category 2021](#), April 2021

<sup>5</sup> *ibid*

<sup>6</sup> HSBC Global Investment Research, *Global Luxury Goods*, March 2026

<sup>7</sup> eMarketer, [US Ecommerce by Category 2021](#), April 2021

<sup>8</sup> Forbes, [Why OpenAI's Checkout Retreat Spells Trouble For Its Commerce Strategy](#), March 2026

<sup>9</sup> PayPal, [PayPal and Perplexity Launch Instant Buy Ahead of Black Friday](#), November 2025

<sup>10</sup> Reuters, [Court temporarily allows Perplexity shopping 'agents' on Amazon](#), March 2026

<sup>11</sup> Bain & Company, [Agentic AI Commerce Hinges on Consumer Trust](#), July 2025

- 
- <sup>12</sup> IAB, [When AI Guides the Shopping Journey](#), October 2025
- <sup>13</sup> Bain & Company, [Agentic AI Commerce Hinges on Consumer Trust](#), July 2025
- <sup>14</sup> Capital One Shopping Research, [Coupon Statistics](#), January 2026
- <sup>15</sup> IMD, [Generative Engine Optimization \(GEO\)](#), November 2025
- <sup>16</sup> Digiday, [ChatGPT is now 20% of Walmart's referral traffic](#), September 25 2025
- <sup>17</sup> CMS Wire, [Cyber Week 2025](#), December 2025
- <sup>18</sup> Noble, [Company Website](#), Accessed April 13 2026
- <sup>19</sup> ECDB, [Gross Merchandise Value of TikTok Shops worldwide in 2024](#), January 2026
- <sup>20</sup> Pixobo, [Company Website](#), Accessed April 8 2026
- <sup>21</sup> HSBC Global Investment Research, Global Luxury Goods, March 2026
- <sup>22</sup> Gorgias, [Company Website](#), Accessed April 9 2026
- <sup>23</sup> Salesforce, [2025 Cyber Week Predictions](#), November 20 2025