

Agentic AI

The Shopper Schism: Competing When AI Agents Become Your Customer

Paul F. Accornero



Image Credit | Sanjoy

Humans still decide what they need, AI agents execute the actual purchase.

When Amazon Makes a \$30 Billion Bet

Late July 2025 turned into an unexpected puzzle. On calls with retail analysts and advertising executives, I kept hearing the same confused reaction: Amazon had pulled its entire product catalog from Google Shopping. For 31 days, the company vanished from search results where it had maintained 30-60% impression share across the United States and international markets.¹

RELATED ARTICLES

Michael Haenlein and Andreas Kaplan, "[A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence](#)," *California Management Review* 61, no. 4 (August 2019): 5–14.

RELATED TOPICS

[Artificial Intelligence](#)

[Sales](#)

[Marketing Strategy](#)

[Digital Platforms](#)

The analysts I spoke with interpreted this as cost-cutting—Amazon saving commission fees to Google. But the timing bothered me. Nine days before the withdrawal, Amazon had rolled out its AI shopping assistant Rufus to all U.S. customers. By fall, Rufus was driving over \$10 billion in additional annual sales, with users who engaged the assistant completing purchases at 60% higher rates.² Why would Amazon pay Google to reach human shoppers when it was building technology to bypass search engines entirely?

Amazon wasn't optimizing for today's customer. It was positioning for tomorrow's. That customer isn't human. It's an algorithm.

I started calling this the Shopper Schism. For over a century, we've operated under a simple assumption: the consumer (the person who uses a product) and the shopper (the person who buys it) are the same individual. That unity is fracturing. Humans still decide what they need—cleaner laundry, new headphones, healthier food. But increasingly, algorithms execute the actual purchase. The consumer and the shopper are splitting into separate actors, and that split changes everything about how commercial competition works.

Understanding What's Actually Changing

Consider laundry detergent. You walk down an aisle, scan options, make a decision based on price, brand, and package design. The entire process happens within your own cognitive framework.

Now imagine telling your phone: "I'm out of laundry detergent. Order what I usually get, but find something better on grass stains since the kids started soccer." An AI agent accesses your purchase history, evaluates alternatives using efficacy data, compares prices across retailers, checks delivery times, and completes the order. You never see product pages. The agent makes product choices; you just set preferences.

This is the Shopper Schism. You remain the consumer—you'll use the detergent, evaluate whether it works. But you're no longer the shopper. An algorithmic intermediary performs that function—the information search, evaluation of alternatives, the transaction.

Why now? AI capabilities have matured. GPT-4 scores 86.4% on graduate-level exam questions across 57 subjects and achieves 80.9 F1 on reading comprehension tasks requiring information synthesis.³ These systems parse specifications, synthesize review data, and reason about trade-offs.

The infrastructure exists too. Amazon’s Product Advertising API, Google’s Merchant API, and similar platforms provide programmatic access to product catalogs, inventory, pricing, and checkout functionality. AI agents query structured databases, processing information far more efficiently than humans scrolling through listings.

And we’re drowning in choice. Amazon lists approximately 600 million products globally.⁴ Research consistently demonstrates that excessive options reduce both decision quality and satisfaction.⁵ Delegating to an agent isn’t just convenient—it’s a rational response to cognitive overload.

How Algorithms Actually Make Choices

I have spent considerable time over the last 3 months watching AI shopping agents evaluate products. What I found: algorithmic shopping represents something qualitatively different from human shopping, not just quantitatively faster.

Start with memory. An agent has perfect recall of every past purchase, every stated preference, every outcome. Human shoppers forget. This creates systematic differences in how humans and agents evaluate alternatives.

Agents also apply decision criteria with relentless consistency. If price matters 40% and quality 60%, those weights remain constant. Human shoppers make different choices based on mood, time pressure, how options are presented. Agents don’t suffer from these cognitive biases—they suffer from different ones.

Here’s what really changes the game: agents are immune to traditional marketing. Aspirational imagery doesn’t register. Emotional appeals don’t work. Brand mystique means nothing to code evaluating JSON files. Recent Columbia Business School research documents this empirically. AI agents systematically favor platform endorsements like “Overall Pick” while avoiding “Sponsored” products. Different agents tend to converge on similar choices, potentially reducing market diversity.⁶

I saw this watching agents evaluate many categories, including wireless noise-cancelling headphones. Brand Alpha provided complete specifications: active noise cancellation at -35dB (certified), battery life 30 hours with ANC on, Bluetooth codec support, driver size, weight, certifications, warranty terms—all verifiable, all structured. Brand Beta led with narrative: “Immersive audio experience for the discerning listener,” sparse specifications, unverified claims, beautiful lifestyle photography.

The agent chose Alpha. Not because it’s objectively superior—I don’t know whether it is. The agent chose Alpha because Alpha is objectively evaluable. Brand Beta’s \$2 million investment in lifestyle photography and influencer partnerships? Functionally invisible to algorithmic assessment.

This creates uncomfortable implications. When the shopper is an algorithm, complete specification tables matter more than beautiful narratives. Verified measurements matter more than emotional connection. Third-party certifications matter more than brand mystique.

The Hard Questions We Need to Ask

Before jumping to strategic responses, we should pause on harder questions without clean answers.

There are agency costs in the economic sense. When you delegate shopping to an AI agent, that agent supposedly acts in your interest. But whose interest exactly? The agent might be trained by a company with conflicting incentives, receive affiliate commissions, or prioritize partners who pay for placement. The opacity of AI decision-making makes it nearly impossible to verify whether agents serve your interests or someone else’s. We’ve recreated the principal-agent problem,⁷ but with algorithmic agents whose decision-making is far less transparent than human intermediaries.

Market concentration becomes more likely. If agents from a handful of companies—Google, OpenAI, Amazon, Apple—mediate most purchases, those platforms accumulate extraordinary power over which products succeed and fail. We’ve seen how Google’s

search algorithm can make or break businesses. Now imagine that power applied not just to information discovery but to actual purchase execution.

There's also what we might lose. Human shopping involves browsing, stumbling upon unexpected products, discovering new brands. Algorithmic shopping optimizes for efficiency and known preferences. Research suggests that pure optimization for previous preferences can create filter bubbles and reduce exposure to novel options.⁸ That matters for both consumer welfare and market innovation.

I don't have good answers to these questions. Neither does anyone else I've spoken with. But they're important enough that we shouldn't ignore them.

Two Arenas, Different Rules

The Shopper Schism creates a fundamentally split competitive landscape. Winning requires strategies for both arenas, even though the rules differ dramatically.

The first arena is the open web, where AI agents compare products across retailers. Google's Shopping agents, OpenAI's Operator, Perplexity's shopping assistant—all evaluate your products alongside every alternative available online.

The currency is structured data:

- Products need comprehensive, machine-readable specifications with proper Schema.org markup
- Detailed attribute tables and standardized categorization become competitive advantages
- Inconsistencies across platforms confuse algorithms and hurt rankings

Objective verification becomes critical. An agent trusts “23% more energy efficient than category average (Energy Star certified)” far more than “eco-friendly design.” Third-party validation carries enormous weight—Consumer Reports ratings, UL certifications, independent test results matter more than brand statements.

The second arena is the walled garden—Amazon’s proprietary ecosystem and other platforms where platform-specific algorithms govern selection.

Key factors include:

- Amazon’s A9 algorithm considers dozens of signals, many hidden, many shifting without notice
- Platform relationships matter: vendor status, FBA participation, advertising spend, responsiveness
- Platform-owned brands have structural advantages
- Pay-to-play dynamics intensify as organic visibility decreases

The strategic challenge is resource allocation. Optimizing for both arenas requires investment in data architecture for the open web, plus platform relationship management for walled gardens. You need different measurement systems, different metrics, organizational capabilities that historically haven’t existed in marketing departments.

The temptation to choose one battlefield is strong, but abandoning either creates vulnerability. Different consumers default to different platforms and agents.

What This Means for Strategy

I’ve spent the past six months studying 30 executives and asking them about adapting to algorithmic commerce. 43% recognize something fundamental is changing but still lacked any formal strategy. Less than 3% have developed coherent responses with formal strategies in place. Over 53% were still primarily focused on human shoppers. What’s emerged are strategic patterns tailored to different competitive positions.

For Major Brands: The Data Fortification Strategy

You have strong brand equity with humans but need to adapt for algorithmic evaluation. The challenge is transforming brand narrative into machine-readable data without losing human resonance.

This requires infrastructure investment:

- Product data management systems where every SKU has complete technical specifications
- Verified performance metrics backed by third-party testing
- Certifications that carry weight with algorithms (Energy Star, USDA Organic, Fair Trade, UL Listing)
- Structured attributes consistent across all platforms

Diversify platform relationships to avoid over-dependence on any single ecosystem. Maintain strong presence across Amazon, Walmart, Target, and direct-to-consumer channels.

Run dual-track marketing. Continue human-focused brand building through emotional connection. But simultaneously invest in algorithm-optimized structured data. You need both.

Procter & Gamble illustrates this with Tide—maintaining emotional brand equity through traditional advertising while publishing comprehensive technical specifications. They ensure Amazon, Walmart, and Target listings have identical, comprehensive structured data. They pursue third-party validation like EPA Safer Choice certification.

For Emerging Brands: The Transparency Arbitrage Strategy

You lack incumbent brand equity, but you can move faster. The opportunity is transparency arbitrage—compete on data transparency and algorithmic optimization before legacy players adapt.

Make transparency your differentiation:

- Publish comprehensive technical data, third-party test results, verified performance claims
- Make your product the most objectively evaluable in the category
- Document everything in machine-readable formats from launch

Maintain direct-to-consumer presence for brand building while optimizing for major platforms. Design products with algorithmic selection in mind from inception. Sprint toward relevant certifications—for emerging brands, third-party validation provides the credibility that brand heritage provides for incumbents.

Sustainable consumer goods brands like Grove Collaborative exemplify this—leading with verified environmental impact data, comprehensive ingredient transparency, and third-party test results.

For Traditional Retailers: The Experience + API Strategy

You own physical touchpoints but must adapt for algorithmic commerce. The challenge is preserving in-store experience advantages while enabling seamless AI agent integration.

Develop robust APIs enabling agents to check inventory in real-time, access pricing information, and transact seamlessly. Invest heavily in fulfillment speed—same-day or next-day delivery becomes table stakes.

Your competitive advantage is aggregating across brands. Provide AI agents with comprehensive cross-brand comparison capability. Become the destination for agents optimizing across manufacturers.

Invest in physical experiences that can't be replicated digitally—expert consultation, try-before-buy, immediate possession. These remain relevant for high-consideration purchases.

Target combines these effectively: robust API infrastructure, same-day delivery through Shipt integration, distinctive in-store experience, and store brands optimized for both human and algorithmic appeal.

Moving Forward: What Amazon Understood

Amazon's withdrawal from Google Shopping in July 2024 wasn't retreat. It was repositioning for commerce's next chapter, where winning means being chosen not by human browsers but by algorithmic agents executing on their behalf.

Companies that will succeed aren't those who choose between human-focused branding and algorithm-optimized data. Winners will build dual capabilities—maintaining brand resonance with human consumers while becoming algorithmically discoverable and favorably evaluated.

This requires uncomfortable organizational changes. Marketing departments need data architecture capabilities. Product managers need to think about machine-readability from the design phase. Brand teams need to work with technical specialists to transform emotional narratives into structured attributes without losing what makes brands meaningful to humans. These represent genuine organizational disruption.

But consider the alternative. As more purchase decisions get mediated by AI agents, brands optimized only for human shoppers will find themselves increasingly invisible to the algorithms that execute transactions. That invisibility, once established, becomes difficult to reverse. The window for making this transition strategically—rather than reactively, under pressure—is closing faster than most companies realize.

In a conversation last month, a CPG marketing executive told me: “We spent twenty years learning to optimize for Google's algorithm. Now we're learning to optimize for Amazon's algorithm, and OpenAI's algorithm, and whoever else builds an agent that shops. It's exhausting. But the alternative is being invisible.”

That exhaustion is real. The organizational effort required is significant. But Amazon made its bet in July 2024, withdrawing from Google Shopping because it was building for algorithmic customers. The rest of us face a choice: continue optimizing product pages and marketing campaigns for human eyes that increasingly aren't seeing them, or rebuild our commercial infrastructure for the agents that are actually executing transactions.

The question isn't whether this shift will happen—it's already happening. The question is whether you're building capabilities for the customer you have today, or the one you'll have tomorrow. And whether you're making that choice strategically, or waiting until competitive pressure makes it for you.

FIVE IMMEDIATE ACTIONS FOR MARKETING LEADERS

If you do nothing else this quarter, start here:

1. AUDIT YOUR DATA COMPLETENESS

Run a gap analysis on your product information across all channels. Most brands I have studied find that 60% to 80% of key attributes are missing or inconsistent. You can't compete algorithmically with incomplete data.

2. PURSUE ONE HIGH-VALUE CERTIFICATION

Identify the most relevant third-party certification for your category (Energy Star, Fair Trade, USDA Organic, UL Listing) and prioritize getting certified. Algorithms weight these heavily.

3. STANDARDIZE CROSS-PLATFORM INFORMATION

Ensure your product specifications are identical on Amazon, Walmart, Target, and your website. Inconsistencies penalize algorithmic rankings. Create a single source of truth for product data.

4. ESTABLISH A DATA QUALITY ROLE

Create organizational accountability for product information accuracy and completeness. This traditionally hasn't lived in marketing, but it needs to now.

5. TEST AI AGENT EVALUATION

Actually watch how AI shopping agents (ChatGPT, Perplexity, Google Gemini, Claude) evaluate your products versus competitors. You'll immediately see what data gaps cost you.

Don't try to boil the ocean. Pick one arena—open web or walled garden—and become excellent there before trying to win both. Most companies spread resources too thin and end up mediocre everywhere.

References

1. Amazon executed a complete withdrawal from Google Shopping ads between July 21-23, 2025, dropping impression share from 30-60% to 0% across all major markets. The company returned to international markets on August 23-25, 2024, but remains absent from the US market as of November 2025. Mike Ryan, “Amazon Has Made a Dramatic International Exit from Product Advertising on Google Shopping,” Smarter Ecommerce, July 25, 2025; Mark Ballard, Tivuiti analysis, July 23, 2025; coverage in Digiday, Search Engine Land, and RetailWire.
2. Amazon rolled out Rufus nationwide on July 12, 2024. By November 2025, Rufus was expected to generate over \$10 billion in annual incremental sales, with users engaging Rufus 60% more likely to complete purchases. Amazon press releases; Adobe Analytics, “2024 Holiday Season Report.”
3. OpenAI, “GPT-4 Technical Report,” March 2023. GPT-4 achieved 86.4% on MMLU (Massive Multitask Language Understanding), a benchmark of graduate-level exam questions across 57 subjects, and 80.9 F1 score on DROP (Discrete Reasoning Over Paragraphs).
4. Amazon lists between 350-600 million products globally as of 2024, with recent estimates suggesting approximately 600 million unique product listings. Red Stag Fulfillment analysis, July 2024; Amazon 2024 10-K Annual Report; Capital One Shopping research, June 2025.
5. Barry Schwartz, *The Paradox of Choice: Why More Is Less* (New York: Ecco, 2004); Sheena S. Iyengar and Mark R. Lepper, “When Choice is Demotivating: Can One Desire Too Much of a Good Thing?” *Journal of Personality and Social Psychology* 79, no. 6 (2000): 995-1006.
6. Research on AI agent shopping behavior from Columbia Business School demonstrates systematic preferences for platform endorsements over sponsored content. Multiple studies from Stanford Human-Centered AI, MIT CSAIL, and industry research labs document these patterns.
7. Michael C. Jensen and William H. Meckling, “Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure,” *Journal of Financial Economics* 3, no. 4 (1976): 305-360.

8. Eli Pariser, *The Filter Bubble: What the Internet Is Hiding from You* (New York: Penguin Press, 2011); Tien T. Nguyen et al., “Exploring the Filter Bubble: The Effect of Using Recommender Systems on Content Diversity,” *Proceedings of the 23rd International Conference on World Wide Web* (2014): 677-686.

Ethical AI Usage Statement:

Generative AI tools were used to assist with drafting, editing, and text refinement during the preparation of this manuscript. The author developed all conceptual frameworks, conducted the field research and industry interviews, and provided the strategic insights presented herein. The author takes full responsibility for the accuracy, integrity, and conclusions of this work. All interpretations, strategic recommendations, and theoretical contributions represent the author’s original analysis based on professional experience and research.



Paul F. Accornero [Follow](#)

Paul F. Accornero is a Senior Executive at the De'Longhi Group and Founder of The AI Praxis. A practitioner-scholar, he researches the "Shopper Schism" — the structural shift from human to algorithmic consumption. He is the author of the upcoming book *The Algorithmic Shopper* on the future of the industry (2027).