

AI Discoverability Risk: A Structural Risk for Retail

Key Takeaways

- 1 AI-driven discoverability is an emerging structural risk in the retail sector.
- 2 After years of optimizing for Google, retailers are losing the narrative to AI assistants.
- 3 When the algorithm decides who gets seen, retailers with limited brand authority, shared assortments, and thin proprietary content face structural pressure.

The Shift: From Clicks to Inclusion

The retail landscape is entering a structural shift in how consumer demand is generated, controlled, and allocated, driven by the emergence of AI tools that select, compare, and ultimately choose products on behalf of consumers. Search engine optimization has long been a core strategic driver of retail demand, but the mechanism by which consumers discover and validate products is now shifting away from SEO and toward AI-mediated recommendations. Consumers increasingly skip the search bar and ask ChatGPT, Gemini, Perplexity, or Google's AI Overviews directly. Rather than returning multiple links, these tools return three or four options, with a recommendation. If a retailer is not in that short answer, the sale is gone before any comparison happens. Adoption is already broad: according to McKinsey, 68% of consumers reported using at least one AI-enabled tool in the past three months; 62% used AI to compare brands, models, prices, or reviews, and 55% used it to learn about a category or product.

This is rendering SEO less effective, pushing retailers to rebalance their investment to stay visible. AI models pull from a much wider set of sources than a traditional search engine, including from Reddit threads, YouTube transcripts, expert reviews, and editorial coverage. Industry research estimates that a retailer's own website accounts for as little as 5% of what an AI model knows about its brand. The remaining 95% is third-party content the retailer cannot directly control. Retailers that are not consistently brought up in the AI agents' recommendation risk a gradual but meaningful erosion in traffic, demand, and revenue due to lack of product visibility. Less mention leads to less traffic, less first-party data, weaker exposure, and less inclusion in the next AI training cycle. It is a flywheel, and once it starts turning the wrong way, it is hard to reverse. Agentic AI, where the agent not only recommends but also selects, configures, and increasingly initiates purchases under the consumer's pre-authorized guidelines, will only accelerate this shift. McKinsey projects up to \$1 trillion in U.S. B2C agentic commerce revenue by 2030, with global projections of \$3 trillion to \$5 trillion.


This has the potential to accelerate divergence across retailers, reinforcing winners while compressing already fragile operators. In other words, AI-driven discoverability is accelerating a

“winner-takes-more” dynamic, while the weaker and often smaller retailers lose relevance at an accelerating pace.

THE STATS

17% of all shoppers, and 47% of Gen Z, now start product searches with ChatGPT or Claude, bypassing Google entirely (WPP / BrandAsset Valuator, World Retail Congress, April 2026). Only ~5% of what AI says about a retailer comes from the retailer’s own website; the other 95% is reviews, editorial, and forums the retailer cannot control (WPP / BAV).

Retailers that spent a decade winning at SEO are now losing the narrative. 89% of retail CEOs say AI-powered search is the shift that worries them most (WPP / BAV). The risk is concentrated in apparel, mid-tier department stores, consumer electronics, and sporting goods — categories where the AI picks the brand, not the store.



Retailers Most Exposed to AI Discoverability Risk

AI discoverability risk lands hardest on retailers built around broad discovery rather than category authority. Marketplaces like Amazon and Walmart get cited because AI can transact against their breadth. Specialty leaders and luxury houses get cited because they own a category answer. Everyone in the middle is at risk, and the divergence is widening. McKinsey estimates the top decile of retailers will capture more than 85% of sector economic profit, up from 73% a decade earlier.

Retailers exposed to this risk typically share three operating attributes:

- **Limited brand authority.** The retailer is not the obvious answer to a category question and lacks the editorial coverage or specialist reputation AI models surface.
- **Undifferentiated assortments.** Inventory overlaps with peers. National brands are commodity inventory for AI shopping, and the AI has no clean reason to pick one store over another.
- **Thin proprietary content.** Limited investment in structured product data (the machine-readable catalog data, such as Google Merchant Center or schema.org markup), owned editorial, expert-led category guides, or media coverage: the inputs AI models weight most heavily.

The shortest test: ask whether the retailer is the answer to a specific shopping question. Costco answers “where to buy in bulk.” REI answers “where to buy outdoor gear.” Sephora answers “where to buy makeup.” Companies like Kohl’s, Express and Big Lots do not answer anything in particular — and that is the issue.

SECTOR	PRIMARY EXPOSURE	RISK
Specialty Apparel	Style and fit queries; AI styling tools pull from competitor catalogs.	HIGH
Mid-tier Department Stores	No category ownership; AI surfaces the brand, not the store.	HIGH
Consumer Electronics	Spec-comparison queries; AI defaults to brands with rich product data.	HIGH
Sporting Goods	Activity-led queries (“best running shoes”); AI cites brands and category specialists.	HIGH
Mid-price Specialty (jewelry, accessories)	Discoverability mediated by editorial and influencer channels AI compresses.	ELEVATED
Home Improvement	Project-planning queries; local data is critical and uneven.	ELEVATED
Grocery & Drug	Repeat-purchase automation; AI list-building bypasses store choice.	MODERATE
Off-Price / Discount	Treasure-hunt model is structurally less AI-mediated.	LOWER
Luxury Houses	Brand authority and editorial coverage protect AI citation share.	LOWER
Mass Merchants & Marketplaces	Walmart, Target, Amazon, Etsy participate in agent protocols.	LOWER

The squeeze is concentrated in the middle. Three categories we view as more exposed to this risk.

- **Mid-tier department stores.** Kohl’s, JCPenney, and Belk. None owns a category in an AI model’s mental map. Asked “where to buy a good winter coat,” AI surfaces Patagonia, Uniqlo, or Amazon.
- **Specialty apparel without a distinct brand claim.** Express, Banana Republic Factory, Chico’s, J.Jill, Talbots. None owns a category AI could name. Asked “where to buy classic women’s basics,” AI surfaces J.Crew, Everlane, Quince, or Amazon.
- **Browse-driven big box.** Big Lots, Five Below, At Home. No digital tailwind, as competitors gain share through AI-mediated discovery.

How Retailers are Responding

The mitigation playbook is still being written, but four moves are emerging:

- **Narrative Optimization (the new SEO).** Structuring content so that AI models can read it, cite it and recommend it. This includes clean product attributes, real-time accuracy, and content built around the questions shoppers ask.
- **Direct platform integration.** Joining OpenAI’s Agentic Commerce Protocol, Google’s Universal Commerce Protocol, or Shopify’s Agentic Storefronts. Engineering cost is real. The cost of staying out is invisibility.
- **Earned media and citation building.** Editorial coverage, expert reviews, third-party authority.
- **First-party AI experiences.** Walmart’s Sparky, Sephora’s beauty advisor, Home Depot’s Magic Apron. Amazon’s Rufus, Macy’s Ask Macy.

The picture is more layered than a single AI path. The largest retailers are not just trying to show up in ChatGPT’s answer, they are also building their own AI layers to sit between the consumer and the generic models. Amazon’s Rufus and Walmart’s Sparky are attempts to keep the recommendation, the comparison, and the transaction inside the retailer’s walls. The strategic question for every other retailer is which channel they are optimizing for: the generic AI that recommends them, the proprietary retailer AI that lists them, or both. The three responses below illustrate how differently retailers are answering that question.

Three Strategies in the AI Game

<p>Build-A-Bear <i>Forced retreat from search.</i></p> <p>CEO Sharon Price John blamed weak online sales on Google’s AI rollout disrupting SEO. Response: cut reliance on organic search, upgrade product schema for AI, lean into direct email and social.</p>	<p>Amazon <i>Building its own AI Moat</i></p> <p>Rufus has 250M+ users with engagement up ~400% YoY. ~\$70B in ad revenue depends on Amazon mediating purchase decisions. Response: block external agents, sue Perplexity, commit \$200B in FY26 capex to AI infrastructure.</p>	<p>Walmart <i>Welcoming every agent</i></p> <p>~4,600 stores within 10 miles of 90% of Americans means margin sits in fulfillment, not discovery. Response: open up to Sparky, ChatGPT Instant Checkout (Oct 2025), Gemini via UCP (Jan 2026). Every agent transaction is incremental.</p>
--	---	---

The Bottom Line

In our February 18, 2026, AI in Retail analysis, we covered how AI is acting as an operational efficiency lever. AI can operate as a discoverability filter at the same time, and the two effects do not net out. Retailers can capture efficiency gains from operational AI deployments in supply chain forecasting, labor scheduling, computer vision, and demand prediction, while simultaneously losing share through AI invisibility, and the second effect outweighs the first for retailers limited in their response.

The winners are retailers with scale, dominant category positions, strong reputational standing, and the engineering infrastructure to participate across multiple AI surfaces. The losers are mid-tier, undifferentiated retailers in apparel, department stores, electronics, and sporting goods, where AI recommendations are largely outside their control and SEO investment no longer carries the weight it did. The key risks we are watching: increased revenue volatility, weaker marketing efficiencies, and structurally higher customer acquisition costs in the exposed segments.



Ramona Bolozan

Director of Analytical Support

email: ramonad@retailstat.com

O: 516. 918. 8855

[LinkedIn](#)

