

bazaarvoice®

The anatomy of an AI-recommended product

A GOOGLE AI VISIBILITY STUDY



Every search is now a storefront

If a shopper asks a Google AI tool, “What’s the best waterproof mascara for someone with sensitive eyes?”, it builds that shopper a personal storefront.

That storefront has three parts. There’s the answer itself, written in plain language. There’s a set of citations, the sources the AI leaned on to write that answer. And increasingly, there’s a product carousel: A set of product cards with star ratings, reviews, and quotes from real customers.

2

The citations

1

The plain-language answer

3

The product carousel showing two to three products with star ratings and reviews

Best waterproof foundation for oily skin

The absolute best waterproof foundation for oily skin is **Dewy Skin Foundation.**

When shopping for oily skin, choosing a silicone-based foundation ensures your makeup stays matte.

Top Waterproof Foundations for Oily Skin



Budge proof foundation products for oily and extremely sweaty skin?

2 Sept 2026 – Some products that may help with oily and sweaty skin include “Powder” ...

5 Best Waterproof Foundations That Stay Put through Sweat...

11 Jul 2025 - Our Top Waterproof Foundations - Best Overall...

The best foundation for oily skin – for lasting coverage that won’t slip

24 Feb 2026 – Which foundation type is best for oily and acne-prone skin?

Show all

Almost every consumer brand is competing for space in that storefront now, and asking: How do I show up in it?

After analyzing over 12,000 Google AI-recommended products across a dozen shopping categories, clear patterns emerge in what the winners have in common.

What earns a spot in the AI storefront

Let's start with the obvious question: Does a strong star rating get your product recommended by AI?


Kind of. AI-recommended products carry a median 4.60-star rating—about half a star above the broader Google Shopping average of 4.37. So yes, your product needs to clear this bar.

But if that's true, why isn't your 4.73-star product consistently beating your competitor's for a spot on the shelf?

Here's the most likely answer: Volume. The median AI-recommended product carries 949 reviews. A typical organic Google Shopping listing carries 120. That's an 8x gap.


Median star rating

Google AI-recommended product vs.
Google Shopping benchmark



Median review volume

Google AI-recommended product vs.
Google Shopping benchmark




To test this, we compared products competing for the same search query head-to-head. Review volume showed a meaningfully stronger relationship to ranking than star rating did. No single signal fully explains what these platforms weigh—they remain, by design, something of a black box—but once a product clears the ratings bar, review volume does more to move it up the AI shelf than having an even higher rating does.

When do product citations appear?

Here's another interesting thing we found: Product carousels aren't evenly distributed across an AI-assisted shopping journey.

They appear in 55.4% of broad, top-of-funnel prompts, like "best matte foundation for oily skin." That drops to 31.1% for mid-funnel comparison prompts, like "Pixel 8 vs. Samsung S24 photo quality," and to 19.5% for bottom-of-funnel prompts, like "best place to buy Pedigree in bulk."

Product carousel frequency by funnel stage in AI-mediated shopping journeys



Where does AI do its shopping?

The citation sources shift throughout the journey too. For example, in response to top-of-funnel prompts, AI leans on content from retailer product pages for citation more often than user-generated content from social media. Mid-funnel, that switches, with UGC from social carrying more weight as shoppers compare specifics.

The result is that different content works harder at different stages of an AI-mediated shopping journey. What earns visibility for a top-of-funnel shopper is very different from what earns it for a mid-funnel or close-to-purchase shopper.


What's overall true of citation sources across the journey though, is that AI tends to lean much more heavily on earned media. These independent, third-party sources account for well over 40% of citations. Owned media—a brand's hosted product and category pages—account for a much smaller share, in the high single digits. If that's because these brand owned pages are not optimized for AI visibility, we don't know from this study. But we do know that Bazaarvoice clients who have made their ratings, reviews, and visual user-generated content visible to AI crawlers using the Authentic Discovery API™, have earned a 40% increase in AI-driven referral traffic.

+40% lift in AI referral traffic

Early adopters have seen a 40% median increase in AI-driven referral traffic after implementing the Authentic Discovery API.



Overall citation sources in AI-mediated shopping journey



All this to say, your [brand.com](#) pages still matter. The content on these pages do a lot of leg work once a shopper lands there. And many shoppers *do* still land there, with 43% preferring to purchase AI-recommended products on the brand's website. It's also worth mentioning that how AI weights different citation sources changes constantly (case in point: the sudden drop in Reddit citations in ChatGPT in late August 2026), and owned media is one of the few fully controllable citation sources for businesses.

The playbook changes by category

If review volume is what wins a spot on the AI shelf, does that mean you need thousands of reviews just to compete? Not necessarily, and certainly not from day one. To begin with, the amount of reviews you need depends heavily on your category:



Health & beauty sits at the top of the scale: a median of 4,700 reviews per AI-recommended product, the deepest of any category we studied. Because competition is so dense in this category, deep review volume may act as a critical consensus signal before Google AI will recommend a product.



Consumer electronics sits much lower, with a median of 1,100 reviews per AI-recommended product. Interestingly, this category is the only one of the 12 studied where star rating carries more weight than review volume for AI ranking. This could be because quality perception is so paramount in this category.



Food & beverage is a distribution story more than a volume story. Retailer product pages are the single most-cited source in this category, and major general retailers dominate where carousel products ultimately point shoppers to buy. Getting reviews in front of those retail partners matters as much as collecting them.




Pet supplies runs on community trust. Discussion forums are the leading citation source—pet owners citing other pet owners—and the retailers AI sends shoppers to are specialists, not general merchandisers.


Category snapshot of top citation sources




Health & beauty




Consumer electronics




Food & beverage



Pet supplies



Category snapshot of median review volume of AI-recommended products



You don't need to be the market leader to compete here

It's easy to look at an 8x review volume gap and assume this game only rewards brands that already have thousands of reviews. That's not what the data shows.

Newer products and smaller brands face an initial headwind—but it's a solvable one, and treating it as a temporary launch challenge rather than a permanent disadvantage changes what's possible. (If you're already ahead on volume, the same logic runs in reverse: protecting that lead means keeping pace with review recency and distribution, not assuming volume alone holds the position.)

Once your product gets a few reviews, a flywheel effect often kicks in: your product shows up higher on the AI shelf, more shoppers discover it, more shoppers buy it, more shoppers review it, and the wheel continues.

A few practical moves:



Prioritize your anchor products. Concentrate review-sourcing effort on the products that matter most, rather than spreading your investment thin across an entire catalog.



Sample strategically before launch. A well-targeted sampling campaign can establish an initial base of authentic reviews before a product goes fully live—giving AI context-rich content from day one.



Distribute what you already have. Existing reviews syndicated to retail and platform partners multiply their reach without requiring a single additional review to be written.



Feed your content directly to Google. Work with user-generated content vendors that provide a direct, structured feed between your content and the Google Shopping Graph (like Bazaarvoice!).

Where to focus next

Four things tend to determine whether a brand shows up well in this new kind of shopping experience:



Source it. Generate customer evidence that's genuinely relevant to how shoppers in your category evaluate products—not just more volume for its own sake. This starts with consistent organic review collection, supplemented by targeted sampling programs where a category or product needs an initial base of authentic reviews to build from.



Amplify it. Extend that evidence beyond your own site into the retailer and syndication network AI already treats as a trusted source. Bazaarvoice Syndication pushes reviews out to retailer partner pages—exactly the kind of destination this research shows AI already leans on.



Make it accessible. Content AI can't access doesn't help you. Much of the web's review content sits behind client-side scripts that AI crawlers can't execute. Serving it in a way crawlers can parse directly—what Bazaarvoice's Authentic Discovery API is built for—has shown a 40% median lift in AI referral traffic across our clients who have turned it on.



Measure it. Track visibility by product, category, and prompt over time. Bazaarvoice partners with independent measurement leaders like Profound, Bluefish, and Evertune rather than scoring its own performance.

None of this promises complete control over a system that remains, at some level, a mystery. But brands that build deep, accessible, and widely distributed customer evidence are consistently the ones AI keeps recommending, and humans on the other end of the AI keep choosing.

About this research

This report draws on Bazaarvoice's internal analysis of 12,745 AI-recommended products across Google AI Mode, spanning 12 consumer shopping categories and three funnel stages (top-, middle-, and bottom-of-funnel), conducted in July 2026. Head-to-head correlation analysis compared products competing within the same search query to avoid any single high-volume product skewing results across the broader dataset. Full methodology available on request.