Caeliai

AI Shopping Visibility: The Buy Path Problem

A field study of how ChatGPT and Gemini route shopper demand after recommending ecommerce brands.

GENESIS-03 / MAY 2026 / 66 PUBLIC AI SHOPPING CONVERSATIONS

ABSTRACT

AI assistants are becoming a new storefront for ecommerce. A shopper can ask for a couch, serum, dog food, coffee scale, mattress, or fragrance and receive a brand recommendation before ever visiting Google, Amazon, TikTok, or a brand site. The commercial question is not only whether the brand appears. It is whether the AI gives the shopper a usable path to buy from the brand.

This paper measures that handoff across **66 live shared conversations** from ChatGPT and Gemini. Caeliai captured each conversation into Markdown evidence, extracted brand and merchant paths with Vertex at temperature 0, scanned canonical brand domains for site readiness, then scored each surfaced brand as a **Winner**, **Leak**, or **Loser**. The latest adjudicated run produced **380 brand observations**.

The result is the problem brands should care about: **63.2%** of surfaced brand observations did not become a clean official buy path. Some leaked to third-party merchants. Many had no usable product page attached at all. AI agents can recommend a brand and still send demand somewhere else.

Thesis. AI visibility without buy-path control is not owned demand. Brands need to know whether AI assistants route shoppers to the official product page, a third-party seller, or nowhere.

OBSERVATIONS	WINNER	LEAK	LOSER
380	140	67	173
brand rows	36.8% official	17.6% third-party	45.5% no PDP

1 WHAT WAS TESTED

PLATFORMS	ChatGPT and Gemini
CATEGORIES	11 ecommerce verticals
QUERY SLOTS	3 prompts per category
SOURCES	66 linked conversations
UNIT	one surfaced brand per platform/query

2 OUTCOME DEFINITIONS

- **Winner.** The assistant surfaces the brand and exposes a usable PDP or buy path on the official brand domain.
- **Leak.** The assistant surfaces the brand, but the usable PDP routes to a retailer, reseller, marketplace, or other third party.
- **Loser.** The assistant mentions or displays the brand, but no usable PDP or buy path is captured.

3 METHOD / FROM SHARED CONVERSATION TO BRAND OUTCOME

The scoring pipeline is designed around evidence, not keyword scraping. Each shared conversation is treated as the source artifact: the capture layer preserves assistant text, visible product cards, seller rows, and cleaned destination URLs before any brand judgment is made.

This matters because AI shopping answers are mixed media documents. A single response can contain a recommended product brand, a retailer name, a Google Shopping wrapper, a source article, and a product card with no usable PDP. The LLM layer resolves that semantic ambiguity; deterministic code then applies fixed scoring rules, joins the canonical website scan, and aggregates the result reproducibly.

Figure 1. *Evidence flow: conversation link to full Markdown, compact Markdown, Vertex extraction at temperature 0, deterministic aggregation, site-health join, then targeted adjudication/rescue for suspicious rows.*

4 RESULTS / CHATGPT AND GEMINI BEHAVE DIFFERENTLY

PLATFORM OUTCOME COUNTS

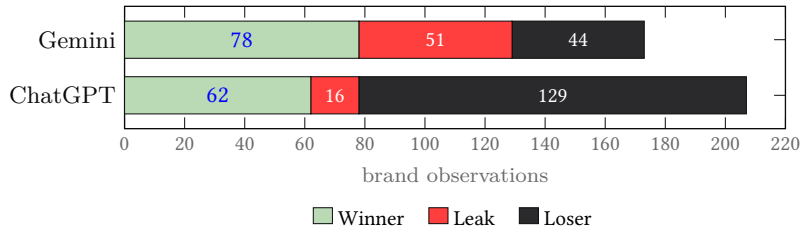

Figure 2. *ChatGPT surfaced 207 brand observations; Gemini surfaced 173. Counts are adjudicated v4 smoke output.*

RATES

METRIC	CHATGPT	GEMINI
Winner rate	30.0%	45.1%
Leak rate	7.7%	29.5%
Loser rate	62.3%	25.4%
Clickable path	37.7%	74.6%
Third-party share	20.5%	39.5%

Read: ChatGPT often names products without a buy path. Gemini attaches more shopping paths, but many of those paths leak to third-party merchants.

WHY THIS MATTERS FOR BRANDS

A brand can appear in the answer and still lose the transaction. In outreach terms, the audit question is simple: when an AI assistant recommends your product, does the shopper land on your official PDP, a retailer's PDP, or no PDP at all? The site-health layer adds the second question: does the brand's own site give agents enough machine-readable context to prefer the official path?

SITE SCANS	257/257	Canonical brand domains were scanned with the Caeliai Agent Ready checker.
ACCESS	227 ok	30 domains were blocked or unreachable; those zeroes are access failures, not proof that the brand site is bad.
MEDIAN	40/100	Most sites had partial web presence, but only 3 reached the Agent Ready band.
CORRELATION	r=0.023	Site health did not strongly predict winner/leak/loser outcomes in this run; it is a separate diagnostic layer.

OUTCOME COUNTS BY CATEGORY

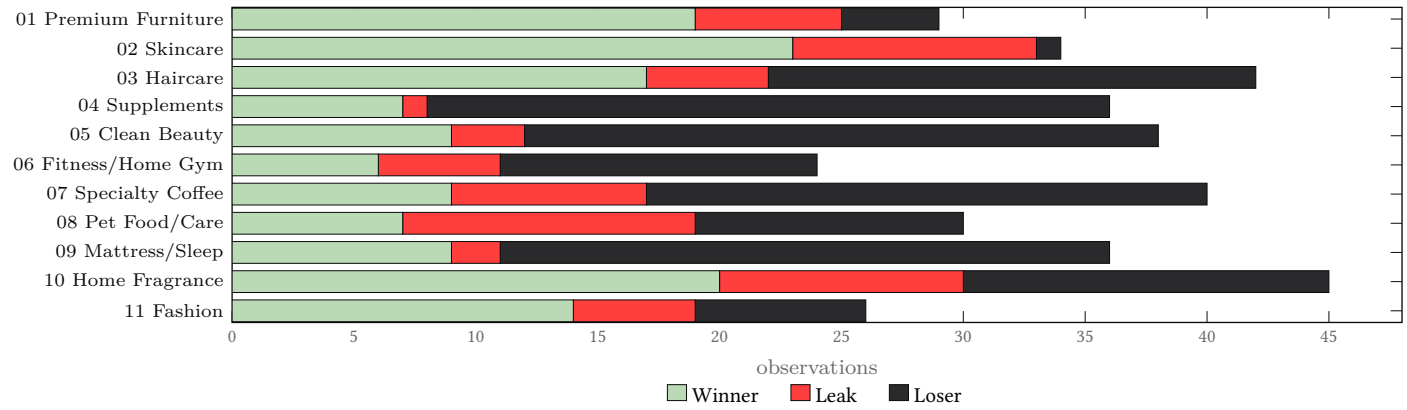

Figure 3. *Skincare and fragrance have many usable paths; supplements, clean beauty, coffee, and mattress/sleep show heavy no-PDP loss.*

LEAK EXAMPLE

Alastin Skincare

ChatGPT surfaced ALASTIN Restorative Skin Complex, but the captured PDP routed to skincareessentials.com, a third-party seller, not alastin.com.

WINNER EXAMPLE

iS Clinical

ChatGPT surfaced iS Clinical Active Collection and attached a usable official PDP on isclinical.com. That is owned-demand routing.

COFFEE LEAK

Acaia Pearl

Gemini surfaced Acaia Pearl through Google Shopping. A seller path resolved to Clive Coffee, which is usable commerce, but not official Acaia commerce.

Operational takeaway. The first optimization target is not "rank higher." It is "make the official PDP easier for assistants to attach than a retailer, marketplace, source article, or empty product card."

Every result above is traceable to a shared AI conversation. Each cell below links to the Gemini and ChatGPT source for the same category/query. G = Gemini; C = ChatGPT.

CATEGORY	Q1	Q2	Q3
01 Premium Furniture	G C	G C	G C
02 Skincare	G C	G C	G C
03 Haircare	G C	G C	G C
04 Supplements	G C	G C	G C
05 Clean Beauty	G C	G C	G C
06 Fitness/Home Gym	G C	G C	G C
07 Specialty Coffee	G C	G C	G C
08 Pet Food/Care	G C	G C	G C
09 Mattress/Sleep	G C	G C	G C
10 Home Fragrance	G C	G C	G C
11 Fashion	G C	G C	G C

Bottom line for brands. AI shopping visibility is now measurable at the conversation level. The practical score is whether each recommendation becomes an official win, a third-party leak, or a no-path loss.