ESSAY

After an AI recommendation, a bid on the brand's own name is defense, not discovery.

By Mike Millett with Digilu

November 15, 2026 · 8 minute read

Panel data on what people search after ChatGPT, Claude or Gemini recommends a brand shows the searches carry the recommended name while unnamed rivals stay flat. The choice is made inside the answer, and the search is where it gets checked.

Marketing teams usually read the traffic after an AI recommendation as a tracking problem: AI influence is invisible, so a new tool must make it visible. The more useful reading is about behavior, and it changes what a bid on the brand's own name is worth. The claim is deliberately narrow: it covers the buyers who act on a recommendation. In the panel data below, more than nine in ten users who saw a brand recommended did not search for it within a week. Every figure here describes the minority who did.

Only 8.8% of visits after an AI recommendation are credited to the answer

Visits that follow an AI brand recommendation are mostly credited to search and direct traffic, not to the answer. Consumers who got a brand recommendation from an AI assistant were 2.5 times more likely to visit that brand's website within seven days, according to Similarweb data cited in the IAB's AI-Era Attribution Blueprint and reported in Digiday's November 12, 2026 media briefing, "The IAB finds AI measurement standards elusive for publishers amid platform data gaps." Only 8.8% of those visits came directly from AI referrals. Search accounted for 55.9% and direct traffic for 19.9%, a combined 75.8%.

Two behaviors could produce that split. Discovery is the work of finding options: a search to learn what a category holds. Verification is the work of checking one choice: is the brand real, is it current, and does its site match what the answer said? The 19.9% direct share leans toward verification, since a

buyer who types an address knows where they are going, though direct traffic also collects bookmarks, returning customers, untagged links and dark social. The IAB figures do not split the 55.9% search share into branded and generic queries. Panel data on what buyers searched after an AI answer does.

Searches after an AI recommendation carry the brand's name, and unnamed rivals stay flat

The searches that follow an AI recommendation carry the name of the brand the answer recommended. In From Prompt to Purchase: How AI Brand Recommendations Move Consumers on the Open Web, a preprint posted on arXiv in June 2026, Michael Iannelli and Alan Ai of Scrunch AI linked an opt-in clickstream panel to the same users' ChatGPT, Claude and Gemini conversations. They then measured what users with no recent observed engagement with a brand did in the seven days after a response, against those users' own earlier weeks.

Discovery predicts a broad effect: the answer sends buyers into the category, brands it did not name gain alongside the one it did, and a neutral mention works about as well as a recommendation, since either hands the buyer a lead. The Scrunch AI panel shows the opposite. Google searches containing the recommended brand's name rose 4.3 percentage points. Unnamed brands from the same category rose 0.6 points. A neutral mention lifted searches 1.8 points: "being named at all does something, but a recommendation does two to three times more," in the authors' words.

A mention alone is worth about 1.2 points over an unnamed rival, and the endorsement adds the remaining 2.5, so most of the lift comes from the recommendation, which happens inside the answer. The Similarweb figures come from a different dataset, so the link is an inference, but read against this panel the 55.9% search share most plausibly records buyers searching the name they were given, not buyers exploring a category.

Searches for a recommended brand grow faster than visits to its site, as verification predicts

Evaluation is a third reading of the search after an AI recommendation: the answer puts a brand on a shortlist, and the search opens a comparison instead of checking a finished choice. Evaluation and verification both predict more searches for the named brand. The rates in the Scrunch AI panel tell them apart.

Before a recommendation, 2.9% of users searched for the brand in a seven-day window and 3.1% visited its site. After one, 7.2% searched and 5.5% visited: a 148% search lift against a 77% visit lift, measured on this panel alone and separate from the Similarweb figure above. Visits outnumbered searches before the recommendation, and searches outnumbered visits after it. Even if every visitor also searched, at least 1.7 of every 100 recommended users searched for the brand and never visited its site that week. Some of them may have looked and rejected the brand. Others had what they

needed from the results page: the brand was real and its listing matched the answer. Either way, the gap is consistent with verification. An evaluating buyer goes to the site, where prices, specifications and terms sit.

The panel cannot see comparisons made inside the conversation, and the brand still arrives at the results page already named, so the search is where a choice is checked, not where it is made.

Last-click attribution credits the verification search and misses the selection

Last-click attribution gives the sale to the final touch it can observe, and after an AI recommendation that final touch is the verification search or the typed address. One path in the IAB report shows what it misses. An answer engine draws on a publisher's review to answer a buyer's question. The buyer reads the answer, never clicks the citation, and buys. The retailer's report shows one line: a sale that "gets reported as a search or direct sale by the retailer." The review shaped the selection, the answer delivered it, and the search recorded it.

Digiday's account of the report names the cost: "That attribution gap can distort marketing investment decisions, with spend going to channels that capture conversions rather than those that influenced them."

Capture is the right word. A paid ad on a brand's own name sits above an organic listing that already holds the page, and it reaches people who have already chosen. A campaign placed at the verification search will look efficient whether or not it changed anything. Its incremental job is narrower than its reported return: it defends a choice against interception, and that defense has a price a test can measure.

The strongest objection: a buyer who has chosen can still be intercepted at the search

The strongest case against reading the search after an AI recommendation as verification is that verification can fail, so money spent at the verification search protects the choice. A buyer who types a brand name can meet a competitor's ad above the organic result, a reseller listing or a comparison page. If the brand does not bid on its own name, a rival can take a buyer the answer had sent to it. A second objection concerns the evidence: the Scrunch AI authors note that a session where the assistant names a brand is also a session with stronger intent toward that brand, so some buyers may have been heading there anyway.

Interception is real, and it does not move the selection into the search. A bid on a brand's own name to block interception is a defensive purchase. Price it as defense, by measuring what is lost when it is paused, and do not credit it with every sale that passes through it.

Selection bias is a fair concern, and none of these sources is a randomized test. The Scrunch AI panel is an observational preprint, it does not see transactions, and its data came from the authors' own company's production system. Within those limits it controls for the obvious confound: it looks only at users with no recent observed engagement with the brand, compares them with their own earlier weeks, and contrasts named brands with unnamed brands from the same category. The recommendation effect holds under those controls. The panel measures the direction of that effect, not the share of all purchases it accounts for.

The cheapest defense against interception is not a bid. Keep prices, availability and claims on the brand's own pages matching what AI answers say about the brand. A mismatch loses the buyer at the verification search whether or not a rival bids.

A paused brand bid in a matched region shows what the bid was worth

A regional pause test prices a bid on a brand's own name. Its unit is the branded arrival index: branded search sessions plus direct sessions, paid and organic together, per region per week. The test runs in four steps.

1. Pause the bid for a fixed period in one region and leave it running in a matched region. National TV, PR and seasonal demand reach both regions, so they cancel out of the comparison.

2. Set each region's index against the same weeks a year earlier, and exclude any week with a TV flight or PR push in one region alone.

3. Measure the population the bid exists to defend: first-time arrivals, the buyers an AI answer sends to a brand they have not used, and the ones a rival's ad can intercept. Where the analytics separate them, run returning visitors as a second comparison instead of mixing the two.

4. Read the paused region against the matched one. If the thesis holds, most paused paid clicks from first-time arrivals reappear as organic and direct, because the answer had already chosen the brand. The share that does not reappear is the bid's real contribution, and a finance lead should fund the bid at the size of that gap, not at its reported return.

When branded search falls everywhere at once, the fix is upstream of the bid

A branded arrival index that falls in the paused region and the matched region together is not a bidding problem. The bid was running in one of them. The decline points upstream, to the inputs AI answers draw on when they choose a brand. The IAB's retailer path shows one at work: a publisher's review shaped the selection before any search happened.

Reviews, press coverage, product documentation and accurate product pages are the inputs a brand can improve. Spending at the verification search can defend a choice. It cannot create one.

> A campaign placed at the verification search will look efficient whether or not it changed anything.

A brand can measure the footprint of AI selection today, in reports it already owns: branded search, direct arrivals, and what a paused bid leaves behind. On Monday, take brand-term return on ad spend out of the growth report, start the regional pause, and size the bid by what the pause shows it kept. Whatever it did not keep belongs upstream, in the reviews, coverage and pages that AI answers read.

Further reading

1. Digiday, Sara Guaglione's report on the IAB AI-Era Attribution Blueprint and the Similarweb data it cites on visits after AI brand recommendations digiday.com/media/media-briefing-the-iab-finds-ai-measurement-standards-elusive-for-publishers-amid-platform-data-gaps
2. Iannelli and Ai, From Prompt to Purchase: How AI Brand Recommendations Move Consumers on the Open Web, an observational panel study of what users do after ChatGPT, Claude and Gemini recommend a brand arxiv.org/abs/2606.10907

marketing helix

marketinghelix.com

A customer behavior model describing how a customer in motion gives a message attention, and how trust, relevance, and timing then shape whether it draws closer, drifts away, or returns.