ESSAY

Answer engines send few customers, and they arrive already decided.

By Mike Millett with Digilu

November 8, 2026 · 9 minute read

Two studies this fall found AI assistant referrals converting above Google organic: 13.8% against 3.02% across 100 sites, 7.18% against 6.08% across 95. The two lines most businesses use to read their brand's pull are where that customer is hardest to find.

Most businesses read the pull of their brand from two lines in analytics: Direct traffic and branded search. Both lines were built for a customer who finds a brand, leaves, and comes back by typing its name. Answer engines produce a different customer, one who compares options inside a conversation the business never sees and clicks only when the comparison is over. The Marketing Helix calls the force behind that click brand gravity: attraction earned by a brand's condition, meaning the state of the brand as the wider public record describes it, rather than attraction bought with spend. This fall's research shows that brand gravity sends few customers, that they arrive ready to buy, and that the usual instruments are poorly placed to find them.

Organic answer placement is earned through condition, and spend buys a place on a list

Answer engines select brands from the record of their condition, and the clearest evidence is an Ahrefs study of 75,000 brands. Visibility in Google's AI Overviews correlated 0.664 with branded web mentions and 0.392 with branded search volume. Mentions are the wider record: what has been written about a brand across the web. Answer visibility follows that record more closely than it follows how often people type the brand's name, so condition, not search demand, is the selector.

The organic answer is where that selector separates earned from bought attraction. On a results list, a paid result sits on the same page as the organic ones, and spend decides which brands appear and in what order. Where an answer engine sells sponsored units beside or inside an answer, those units work the same way: a bought placement the customer still has to weigh. The organic answer itself is drawn from the record the engine reads, and no bid places a brand in it. Whatever a brand buys, it buys a place in a comparison the customer still has to make.

The distinction fails on one measurable condition, and the paid-versus-assistant test measures it. If bought attraction reached customers who had already decided, paid search visitors would convert like assistant referrals. Any advertiser can run the test this week: put the paid search conversion rate beside the conversion rate of referrals from ChatGPT and other assistants, over the same period and on the same conversion event. Where the two rates match, the distinction between earned and bought attraction does not hold for that business. Where paid converts below the assistant referrals, it does.

Settled arrivals from ChatGPT convert at 13.8%, against 3.02% from Google

Marie Haynes's October 21 analysis of 100 sites and more than 145 million sessions over 12 months found ChatGPT visitors converting at 13.8%, against 3.02% for Google organic visitors. ChatGPT sent 0.62% of the traffic Google organic sent. The Marketing Helix calls the ChatGPT visitor a settled arrival: a visit made after an answer has already narrowed the options. The click records the end of a decision, not its start.

Both groups in the Haynes data are unpaid, so the gap does not compare earned attraction with bought attraction. It compares an answer with a list. An answer narrows the options before the click, and the visitor arrives decided. A list leaves the comparison open, and the visitor arrives still deciding. The 13.8% against 3.02% is the measured distance between those two states.

The bought half of the thesis is a prediction this data does not test. Paid search is a list with a price on each position, so the settled arrival mechanism predicts that paid search converts at or below Google organic, and below assistant referrals. None of the four datasets published this fall reports a paid search conversion rate. The paid-versus-assistant test, run in one advertiser's own account over the same period and on the same conversion event, is where a reader proves the prediction or breaks it.

Two reasons to distrust the 13.8% figure

Marie Haynes's 13.8% conversion rate for ChatGPT visitors has two weaknesses a skeptic should raise before accepting it. The first is pooling. The figure combines 100 sites, and a few sites that convert unusually well can carry a pooled rate. A median or a site-by-site distribution would show whether that happened here. The closest check available is the SEO Works study of 95 GA4 properties. It found AI referrals converting at 7.18% against 6.08% for organic overall, a far narrower gap, and converting worse than organic on 23 of the 63 sites with enough volume to compare. The premium exists in aggregate and is absent on more than a third of the sites that can be measured.

The second weakness is tagging. In October, Search Engine Journal reported that the share of ChatGPT answers with at least one tagged source link fell from about 98% on October 8 to 66% by October 14. Adobe's Cedric Huesler said most ChatGPT clicks arrive without a referrer. A visit without a referrer is not counted as ChatGPT, so Haynes's ChatGPT group is the subset that arrived tagged. The 0.62%

volume figure is therefore a floor, since untagged ChatGPT visits are counted somewhere else. The 13.8% rate holds only if untagged visitors convert like tagged ones, and nothing in this fall's data tests that.

What survives both checks is the pattern rather than the precise figure: a conversion rate above organic in aggregate, uneven across sites, from a volume that stays small even as a floor.

Direct traffic now records bots and misfiled answer engine clicks

Direct traffic was once a rough count of people who knew a brand well enough to type its address. Kevin Indig's October 19 analysis of 12 GA4 properties shows how little of that count remains. Total Direct traffic across his sample rose 140.9% between January and September, while engaged Direct fell 1.0%. On one site, GA4 recorded 409,000 Direct sessions, and 198 were engaged. Greg Jarboe, who compiled the research in Search Engine Journal on November 6, reads that gap as bots, not buyers.

Direct traffic also absorbs real answer engine visitors under the wrong label. Adobe's Cedric Huesler found that Direct rose on two of every three pages that lost ChatGPT traffic. Huesler said a fix began rolling out on October 15.

Indig then tested whether weekly AI visibility predicted later Direct traffic or branded search, running 632 correlation tests on four brands at lags of zero to six weeks. When he reran the tests with deliberately wrong dates, several of the mismatched alignments correlated as strongly as the real ones. His conclusion was that an effect could exist but he could not isolate it. A null result on four brands does not show that answer engines have no influence. It shows that this instrument cannot separate a real effect from coincidence at that scale, and a business reading its own Direct line has one brand, not four.

Branded search counts the customers who already knew the name

Branded search has a better claim than Direct traffic to measure a brand's pull, because a person who types a brand's name is plainly a person. In the Ahrefs study of 75,000 brands cited above, branded search volume tracked AI Overview visibility less closely than branded web mentions did. A correlation across that many brands describes which brands have both, and large, well-known brands tend to. It does not show that one brand's visibility produces later searches for its name, the question Kevin Indig's lagged tests asked and could not answer.

The settled arrival mechanism predicts why branded search lags. A customer who has just seen a brand named in an answer has a link in front of them and no reason to search for the name. Adobe's data shows where that click lands: on two of every three pages that lost ChatGPT traffic, Direct rose. The answer engine visitor reappeared in Direct, unnamed. Branded search grows when a name is remembered and looked up later. It records the customers who already knew the name, not the ones an answer just gave it to.

The case for ignoring answer engines this year

The strongest case against the settled arrival argument comes from the same datasets. Google drove 96.5% of conversions in Haynes's sample, 27 times more than all large language models combined. The SEO Works study found that a 5% lift in organic conversions added 10.3 conversions per site per month, while doubling AI referrals added 6.6. For B2B product businesses, AI referrals converted at 0.62 times the organic rate. On this reading, a business with limited attention should spend it where 96.5% of its outcomes already come from, and the conversion premium is a curiosity of small samples.

That case is right about volume, and the argument here concedes it. The claim concerns measurement, not size. Referral volume from answer engines has been measured, and it is small. The influence beyond the referral is the part that Direct traffic and branded search are poorly built to show. A business can accept the 96.5% figure and still conclude that it is reading its pull from the wrong lines.

The B2B result needs a straight answer. Brand gravity predicts that the premium appears where an answer can settle the comparison and fades where it cannot, as in a purchase that several stakeholders must approve after any answer is given. The B2B product result fits that prediction, but one category is not a pattern. A business can test it on its own data. If its AI referrals underperform on purchases that need approval after the answer and outperform on purchases one person settles, the prediction holds. If the underperformance falls on simple purchases, it fails.

Ask about the decision, not the discovery

Customer surveys reach the one part of an answer engine purchase that analytics cannot: the conversation in which three options became one. Analytics records where a customer arrived. With an answer engine, the comparison ended before that.

Jarboe's column recommends adding a "How did you hear about us?" field and including AI assistants in brand lift surveys. It notes that this data has flaws of its own: memory is selective, and people overweight the most recent or most vivid touchpoint. Those flaws are bounded in a way the Direct line is not. A survey answer comes from a customer who bought, so a faulty memory moves credit between real channels and cannot add a buyer who does not exist. Direct traffic has no such floor: on the site where GA4 recorded 409,000 Direct sessions, 198 were engaged, and nothing caps how far automated visits can inflate the rest.

The discovery question is the weaker form. A customer may have first heard of a brand years ago and decided last week, after an answer engine confirmed what the record said. "How did you hear about us?" captures the first event. The decision lives in the second. Ask instead what else the customer considered, what settled the choice, and whether an AI assistant was part of it. Ask at intake and again at the next purchase or renewal. The record an answer engine reads changes, and the same question asked twice shows whether settled arrivals are growing or fading.

> The click records the end of a decision, not its start.

The four datasets published this fall do not settle how much answer engines matter; they show where to look. An answer engine makes its choice from the record of a brand's condition, in a conversation no analytics line can see, and the customer it sends arrives with the choice already made. Direct traffic and branded search will not find that customer. A question put to that customer will.

Further reading

1. Greg Jarboe, Search Engine Journal: compiles the Haynes, Indig, SEO Works and Ahrefs datasets on AI visibility, conversion rates and Direct traffic searchenginejournal.com/ai-visibility-isnt-a-sham-most-ways-of-measuring-it-are
2. Matt G. Southern, Search Engine Journal: reports the October drop in tagged ChatGPT links and Adobe data showing lost ChatGPT visits reappearing as Direct searchenginejournal.com/chatgpt-traffic-drop-appears-to-be-a-tracking-problem
3. Cloudflare Blog, The Internet has a second audience: network data showing non-human traffic passing half of the internet traffic Cloudflare sees, and AI agent growth blog.cloudflare.com/agentic-web

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.