

SIGNLLABS.

WHITE PAPER

THE ANSWER ECONOMY

How AI decides which brands to recommend, and how to be the one it names.

82–84%

OF AI CITATIONS TRACE TO
EARNED AND THIRD-PARTY
MEDIA

3×

CONVERSION RATE OF AI-
REFERRED VISITORS ·
MICROSOFT, 2026

–33%

GOOGLE REFERRALS TO
PUBLISHERS, YEAR TO NOV
2025

2.5B

PEOPLE SEE AI-GENERATED
ANSWERS EVERY MONTH

By Erika Hanafin Austria, SignlLabs, the AI Signaling practice of GCW. 2026. The systems behind these rankings update monthly, so every figure here is timestamped.

EXECUTIVE SUMMARY

When a buyer asks AI who or what brand is best, the answer engine does not walk them to your website. It just answers with what it knows. It names a few brands it trusts, and it moves on. SignlLabs was built to make sure your brand is the one named. This paper covers three things we know, right now: how the five major answer engines decide,

what earns a citation, and what happens after the click – where the best traffic on the internet is quietly bleeding out. One truth runs under all of it: **your website is a bit player, earned trust is the whole story**, and the buyers AI sends you behave nothing like the ones you built your site for.

01 THE SHIFT: FROM SEARCH TO ANSWERS

Buyers are not Googling the way they used to. They ask ChatGPT a question and get a curated answer in seconds. No links clicked, no landing page visited, no form filled. Most brands still measure impressions and visibility as if buyers are walking through the front door. That front door

barely exists anymore. For a growing list of categories, there is just an answer, generated once and read by someone who never scrolls further.

This is not a forecast. It is already here, it is fast, and it is quietly moving your most valuable traffic somewhere else. The numbers, as of 2026:

SCALE

2.5B / clicks halve

Google's AI Overviews alone now reach roughly 2.5 billion people every month ^[1] – and when one appears above the results, clicks to the sites below it roughly halve ^[2].

THE NUMBER THAT SHOULD REORDER YOUR BUDGET

3x conversion

Visitors AI refers convert at roughly three times the rate of everyone else (Microsoft, 2026). AI does not send more traffic. It sends better traffic. Buyers who already did the comparison and showed up ready to say yes.

THE REFERRAL COLLAPSE

–33%

Google sent publishers about 33 percent fewer visitors in the year to November 2025. The answer absorbed the click.

A MARKET FORMING IN REAL TIME

\$365M / +40% CAGR

The US generative engine optimization market sits near 365 million dollars in 2026 and is compounding above 40 percent a year. The language of this category is being written now.

Read that again. The brand winning the ad auction is often not the brand the machine recommends. If you are missing from the answer, you are not just losing a click — you are losing the buyer who already compared, already decided, and was one recommendation away from choosing your brand.

02 HOW ANSWER ENGINES ACTUALLY RESPOND

An answer is built from three layers. You feed each one differently, or you do not show up, plain and simple.

What the answer engine memorized

TRAINING DATA • SLOW, STABLE

Training data. Slow, stable, driven by who you are. You earn your way in by being consistent and structured everywhere the model learns.

What the answer engine looks up

LIVE SEARCH • MOVABLE IN WEEKS

Live search, at the moment the question is asked. Fast, and movable in weeks. Fresh content helps. The pages other people wrote about you (third-party endorsements) help far more.

What the answer engine trusts

EARNED CONSENSUS • LASTS

Facts repeated across many credible, independent sources. The hardest layer to fake and the one that lasts. You do not publish your way here. You earn it.

The mechanic to understand: query fan-out

A prompt is not a keyword. Ask an answer engine a question and it does not answer that question. It quietly breaks your one prompt into about five, researches each separately, then hands back a single answer stitched together from the sources it trusts across all of them (Microsoft, 2026). You cannot keyword your way into five questions at once. You earn your way in with presence that is broad, consistent, and credible everywhere the answer engine looks.

1 → 5

ONE PROMPT BECOMES
ABOUT FIVE
RESEARCHED QUESTIONS

Answer engine by answer engine, mid-2026

ChatGPT, Perplexity, Gemini, and the rest do not behave the same, and their citations barely overlap. One study found only about **11 percent of domains were cited by both ChatGPT and Perplexity**. Winning one is not winning all. Score each answer engine on its own and never trust a blended number.

ANSWER ENGINE	SOURCES	HOW IT RESPONDS (MID-2026)	WHAT WINS IT
ChatGPT	~4	Periodically updated training data plus a search partner and licensed publishers. About 4 sources per answer. Heavy on Wikipedia (roughly 26 to 48 percent of top citations) and Reddit.	Wikipedia and Wikidata presence, Reddit discussion, Bing-indexed pages
Perplexity	~12	Real-time web search and rerank on every query. The heaviest citer, around 12 sources, and the most recency-hungry.	Fresh, retrievable content and third-party coverage; the fast-win answer engine
Claude	2–3	Cites the documents handed to it; rewards long-form analytical writing from named authors. Usually 2 to 3 sources.	Structured, authored, in-depth writing with real expertise
Gemini / AI Overviews	76%	Reuses Google's ranking and index, plus YouTube and Reddit. In mid-2025, about 76 percent of AI Overview citations came from top-10 organic results.	Strong organic footprint, multi-format content, YouTube, Reddit
Copilot	fewest	Bing index, fewest citations, winner-take-most. One definitive page can capture the answer.	One canonical, authoritative page in the Bing index

Here is the part that should stop you. Reddit is the most-cited source across every major answer engine, at nearly 40 percent. Wikipedia is second. Neither one is your website. **The answer engines are mimicking what other people are already saying about your brand.**

~40%

REDDIT SHARE OF CITATIONS, ALL ANSWER ENGINES

11%

DOMAIN OVERLAP, CHATGPT AND PERPLEXITY

26–48%

WIKIPEDIA SHARE OF CHATGPT TOP CITATIONS

5

ANSWER ENGINES TO SCORE SEPARATELY, NEVER BLENDED

03 WHAT ACTUALLY EARNS A CITATION

Four things decide whether the machine says your name.

01 · STRONGEST PREDICTOR

Referring domains

82–84%

How many credible, independent sites mention you? This is the single strongest predictor of a citation, and it lives outside your website. Across 2026 studies, roughly 82 to 84 percent of AI citations trace to earned and third-party media. Your own pages account for just 5 to 10 percent. That gap is the whole game.

02 · FRESHNESS

Recency

6.0 vs 3.6

Pages updated in the last three months average 6 citations. Stale ones average 3.6. Perplexity punishes old content hardest of all.

03 · LIFTABILITY

Structure

40–60 words

Clean 40- to 60-word answers, real question-and-answer blocks, a heading order that follows how people actually ask. Good structure does not win the citation. Bad structure loses it.

04 · IDENTITY

Entity clarity

One name

One name, one description, one set of facts, everywhere, backed by Wikipedia and Wikidata, so the answer engine knows exactly who you are and never hands your recommendation to a brand with a similar name.

Why AI trusts people more than brands

There is a reason the sources that win are rarely your own pages. A brand account is supposed to say nice things about itself, ie: advertising, and everyone knows it, humans and

answer engines alike. An individual is not carrying that bias. When a person says something specific, honest, and a little against the grain of what you would expect a company

to say, it reads as real, or as a testimonial. It is the same reason you trust a friend's recommendation over an ad. The answer engines did not invent that dynamic. They just started keeping score of it.

So here is the implication, and it is not comfortable. Your recommendation has to be earned in other people's voices. In the communities, the reviews, the expert commentary, the coverage the answer engine already trusts. **Your homepage cannot vouch for you. A credible third party can. Build for that, or stay invisible.**

04 SEO, AEO, AND AI SIGNALING ARE NOT THE SAME THING

People use these like synonyms. They are not, and the confusion is expensive. Each one targets a different moment, and only one of them wins the recommendation.

DISCIPLINE	WHAT IT WINS	WHAT IT OPTIMIZES
SEO	A ranking on a page of links	Keywords and backlinks, for a human to click a result most buyers now skip
AEO	The snippet an AI reads back	Structured, liftable answers so a model can quote you. Necessary, but it only makes you eligible
AI Signaling	The recommendation itself	The earned, third-party trust the answer engine leans on, so when someone asks AI who is best, you are the answer

SEO earns the link. AEO earns the snippet. AI Signaling earns the recommendation. Because of query fan-out and the earned-media math, you cannot keyword your way in. You earn your way in, across every source the machine reads.

05 THE HALF NOBODY OPTIMIZES: AFTER THE CLICK

Winning the citation is half the job. The other half is what happens when they land. This is the part almost nobody optimizes, and it is where the best traffic on the internet goes to die. Two truths change how you build.

An AI visitor is not a social visitor

They are opposites. The AI visitor shows up with the comparison already done. Short journey, sharp focus. What they need is confirmation. Price, stock, shipping, proof. The social visitor still needs convincing. Optimize both the

same way and you leave money on the table (Microsoft Clarity, 2026). Split your analytics by source before you judge a single page.

The three post-click killers


Behavioral data surfaces three patterns that kill conversion after a high-intent buyer lands. None of them is a traffic problem. All three are fixed on the page.

Dead clicks in checkout A broken control. The visitor clicks a button, nothing happens, and they leave. Your best buyer, lost to a bug.	Unclear pricing “From \$49,” with the real number buried behind a configurator. Rage clicks, then bounce.	The offer below the fold Free shipping and the primary button, never seen, because most people stop scrolling around the 20 percent mark on every device.
---	---	---

A case study, and the wrong fix

A real pilot brand was sure non-domestic European shoppers were bailing over price or shipping distrust. The data said no. 54 percent of non-domestic visitors landed straight on English blog and research articles, and 96 percent of them bounced, because they were readers, not shoppers. Cart-reach told the story. Domestic visitors 5.63 to 7.11 percent, non-domestic 2.26 to 2.86. Then the tool flagged something the brand never asked about. The English storefront was only half translated. The cookie-consent and Add-to-Cart buttons still rendered in the original language, and bounce rate ran 8.6 percent higher because of it (Microsoft Clarity, 2026). A full pricing overhaul would have been the wrong fix. The right fix was upstream. Better targeting, and a page that actually spoke the visitor’s language.

PILOT BRAND · THE DIAGNOSIS



The lesson for AI Signaling is simple. Measure the whole journey. A citation you win is worthless if the page you win them to cannot close a buyer who arrived ready to buy.

06 THE SIGNLLABS METHOD

Measuring AI visibility is becoming a commodity. Soon everyone will have a number. A number does not move anything. What moves the answer is a method, and the earned-media muscle to run it. That is what SignlLabs is built to be, inside GCW, a public relations firm with 40+ years of experience earning trust for brands across top-tier traditional and online media. That is how recommendable your brand is, scored with the Maturity Method, across six dimensions, five answer engines, and ten real buyer intent questions, and it is re-scored every quarter.



The six dimensions

DIMENSION	WEIGHT	WHAT IT MEASURES
Foundation	24%	Structured identity across owned and earned media. Can the machine tell exactly who you are?
Endorsement	20%	Third-party authority. Reviews, references, credentialed mentions in the sources that count.
Alignment	16%	Do your claims and what the sources say match? One story, everywhere.
Presence	16%	How densely you show up across the sources AI trusts and pulls from.
Warmth	12%	How you are talked about. Conviction, or caveat.
Constancy	12%	How fresh and recent your signals are, every time the answer engine looks.

Endorsement and Presence carry the most weight, because that is where the citations actually live. The score runs zero to one hundred, up a ladder: Unseen, Surfacing, Standing, Trusted, and The Answer.



How SignlLabs moves it, in two layers

LAYER ONE · TABLE STAKES

The owned foundation 5–10%

Ship the schema. Write the answer blocks. Publish an llms.txt and open the AI crawlers. Verify the domain in Bing Webmaster Tools and Google Search Console. Refresh the cornerstone pages every quarter. This is table stakes. It gets you eligible. It does not get you named.

LAYER TWO · WHERE THE SCORE MOVES

The earned authority 82–84%

This is where the score moves. Rank the outlets by how often AI cites them in your category, not by human reach, and pitch there first. Show up, for real, in the communities and review platforms the answer engines quote. Then earn the entity anchors: Wikipedia and Wikidata. Seed original data so trusted sources carry your facts into the index. This is what a 40-year earned-media practice was built to do, pointed at the answer engines. And remember: each answer engine weighs these sources differently, so the approach has to be tailored answer engine by answer engine.

The ten buyer questions

Because of fan-out, you are not scored on one question. You are scored across the whole journey.

01 Discovery and Category	02 Experience and Lifestyle	03 Cost and Value	04 Fit and Qualification	05 Decision Criteria
06 Program and Differentiation	07 Trust, Safety and Reputation	08 Competitive and Comparison	09 Urgency and Trigger	10 Founder and Authority

Ten different questions a real buyer asks an answer engine, ten different signals to earn.

07 PROOF: WHAT THE METHOD MOVES

Here is what it looks like when it works. A national women's health and public-health advocacy brand needed to be the source answer engines cite. We measured it, engineered two data-driven releases through national wire, and tracked the citation signal across all five answer engines, before and after.

7→47 AI CITATION SCORE, OUT OF A POSSIBLE 60, ACROSS FIVE ANSWER ENGINES	+572% CUMULATIVE AI CITATION SIGNAL, IN 18 DAYS	250M+ COMBINED EARNED REACH ACROSS TWO RELEASES	2 of 5 ANSWER ENGINES REACHED A PERFECT CITATION SCORE
---	--	--	---

The mechanism was the wire-to-AI flywheel. The earned syndication taught the answer engines who the trusted source was, and the citations followed in days, not quarters. One answer engine dipped while three surged. We caught

it, because we re-score every quarter, and we report per answer engine, never blended. A blended average would have hidden the dip and cost the client share.

08 HOW TO MEASURE IT

Rank is a vanity number now. Measure what actually happens, on both sides of the click.

TRACK	WHAT TO WATCH
Share of authority	The citations you earn on your topics, branded versus non-branded, week over week. Are you getting cited on questions your competitors used to own?
Per-answer engine citation	Named or not, cited or a passing mention, position, sentiment, on each of the five answer engines. On their own. Never blended.
Bot activity	Which AI crawlers hit which pages, how often, and which they skip. Connect a CDN to see it.
Behavior by source	Split AI from paid, organic, social, and direct, and watch how each one converts. Free in Microsoft Clarity.

A tool worth knowing: Microsoft Clarity is free, and it does both jobs at once. Upstream, it shows bot activity, the queries you are cited for, and your share of authority. Downstream, it shows heatmaps, session recordings, and behavior by source. Verify your domain in Bing Webmaster Tools or Google Search Console to switch citations on. It is a strong baseline to sit next to a managed Maturity Method program.

09 A TWO-WEEK BASELINE

WINDOW	AI VISIBILITY (UPSTREAM)	AFTER THE CLICK (DOWNSTREAM)
Weeks 1–2	Set up topics and prompts, baseline your citations and share of authority, per answer engine	Find the biggest friction points on your top landing pages
Ongoing measurement	Track share of authority week over week. Are you getting cited on topics you were missing?	Confirm the fixes landed. Did cart adds or bookings actually move?

Lock the baseline before your peak season. You cannot prove lift without a clean before. And you cannot fix a page you have never watched a real buyer struggle on.

10 THE LIMITS, PLAINLY

This field is young. The answer engines change every month. So, take every number here as a snapshot with a date on it, not a law. Citation shares move. Anyone who promises you a ranking is guessing or selling. What does not change is the shape of the work. Measure per answer engine, and build the owned foundation to earn the

authority. Fix the page the buyer actually lands on. Re-score on a cadence, so drift gets caught instead of discovered. **The operators who date their claims and show the before-and-after are the ones who will win this decade of visibility.**

ABOUT THE AUTHOR

Erika Hanafin Austria leads SignlLabs, the AI Signaling practice inside GCW, a public relations firm that has been earning trust for brands for 40+ years. We help brands get named by the answer engines. Measured with their proprietary Maturity Method, and enhanced with earned media. Find me at signllabs.ai.

Cite this paper: Hanafin Austria, Erika. "The Answer Economy: How AI Decides Which Brands to Recommend." SignlLabs, 2026. signllabs.ai.

SOURCES

1. Microsoft Clarity, AI visibility and behavioral analytics briefing (Microsoft Advertising, 2026)
2. AI Citation Patterns Explained (GrackerAI, 2026)
3. How each platform cites sources differently (Discovered Labs, 2026)
4. AI Platform Citation Source Index 2026 (5W / PR Newswire)
5. Generative Pulse: Earned Media Drives AI Citations, ~84% (GlobeNewswire, 2026)
6. Top sources LLMs cite most (Contently, 2026)
7. Mastering Generative Engine Optimization in 2026 (Search Engine Land)

REFERENCES

- [1] Google Says AI Overviews Has Over 2.5 Billion Monthly Active Users as New Website Controls Roll Out – Digital Information World, June 3, 2026
- [2] Google AI Overviews Slash Click-Through Rates Almost in Half – TechNow