

12,792

AI ANSWERS ANALYZED

AI Is Citing the Answers Businesses Publish

What 12,792 AI answers and 3,127 business websites reveal about local recommendations, reviews and buyer questions.

3,127

business websites

8

service industries

12

US metros

AI Visibility Study · September 2026 · By Liron Segev

AnswerContentEngine.com

WHY WE DID THIS

When someone asks AI about a service, whose information becomes the answer?


We tested which businesses 3 AI models named, compared their websites with businesses the models never named, and examined the pages cited when buyers asked about costs and choices.

The mission was to understand what actually works to get AI to recommend your business, and what is just hype.

KEY FINDINGS

7 of 96

More reviews didn't win. The review leader was AI's top pick in just 7 of 96 markets.



■ Review leader was also AI's most-named business (7) ■ Did not match (89)

We checked 96 local markets. Each market is one service in one city, such as plumbers in Denver. In each one, we compared the business with the most Google reviews against the business AI named most often.

They matched in only 7. In the other 89, the review leader was not the business AI named most.

Reviews still matter. But having the most reviews rarely made a business the one AI named most often.

AI's existing knowledge, with web search turned off.

KEY FINDING 02

99 vs 44


The businesses AI named had more than twice as much website content.

■ Typical business AI named ■ Typical business AI never named

Website pages



Pages answering buyer questions or comparing options



The typical business AI named had 99 pages on its website. The typical business AI never named had 44. The named businesses also had 5 pages answering buyer questions or comparing options. The others had 1.

“Typical” means the middle business in each group. We counted pages, not words.

KEY FINDING 03

59.6%

AI named the business with more content in 6 of 10 matchups with the same service, same city and similar reviews.

OUT OF EVERY 10 MATCHUPS, THE BUSINESS AI NAMED HAD MORE PAGES IN 6



To compare like with like, we paired businesses head to head: same service, same city, similar number of Google reviews. In 59.6% of those pairs, the business AI named had more website pages.

That is not a coin flip. The difference is statistically significant.

348 pairs, 4 ties excluded. Clearest when AI answered from existing knowledge; the web search comparison was inconclusive.

KEY FINDING 04

The buyer's question decides which page AI links to.

ASK WHO TO HIRE**57.9%**

of the business links AI referenced were homepages.

OUT OF EVERY 10 LINKS**ASK WHAT IT COSTS OR WHICH OPTION TO CHOOSE****70.6%**

of the local-business pages AI linked to were pages with answers or pricing. That is 7 out of every 10 pages: articles, FAQs and pages explaining prices.

OUT OF EVERY 10 PAGES


In plain terms: when buyers want a name, AI points to your front door. When they want to understand costs and choices, AI points to the page that explains them. If that page does not exist on your website, AI cannot cite it.

Different question sets and samples: 739 links to business websites on hiring questions, 177 local-business pages on cost and comparison questions.

KEY FINDING 05

36.0%

Local businesses' own websites were the biggest source AI cited on cost and comparison questions.



More than 1 in 3 pages AI cited came from local businesses' own websites (36.0%). That was more than cost-guide sites, national firms and blogs combined (32.3%), and more than 5 times the share from directories (6.7%).

Your website can be a source AI uses to answer your buyers' questions.

288 answers to cost and comparison questions with web search on; 492 page citations.

KEY FINDING 06

55%

55% of the local businesses AI linked to were not in the first 60 Google Maps results.

THE 129 LOCAL BUSINESSES AI LINKED TO



- 71 not in the first 60 Google Maps results
- 58 in the first 60

When buyers asked about costs and choices, AI linked to the websites of 129 local businesses. 71 of them did not appear in the first 60 Google Maps results we collected for their service and city.

Not leading the map did not stop their pages from being cited.

Maps results depend on the query, location and collection time. “Local” includes businesses elsewhere in the same state.

These findings measure different things: website characteristics, business mentions and page citations. The following pages explain the sample and meaning of each.

THE INVESTIGATION

Here are the details of what we did:

We didn't just ask AI a handful of questions. We collected **12,792** answers from OpenAI, Claude and Gemini models. Our core local-business research covered **12** metros and **8** service industries.

Then we examined **3,127** business websites to investigate what separated the businesses AI named from those it never mentioned.

We built the investigation around **96 local markets**, repeated fixed questions, and tested AI both with and without web search.

We checked business identities against Google Maps, compared website content, and matched businesses with similar review counts to test whether the content difference remained.

A separate experiment examined the sources AI cited when buyers asked about costs and choices.

That depth matters.

An isolated answer can reveal a business name.

Repeated, structured testing lets us examine how consistently businesses appear, compare them with their competitors, and assess whether an observed difference is statistically significant. **The findings in this report come from that investigation.**

Our investigation had 3 stages

1

Ask which businesses AI names

12,504**AI ANSWERS**

3 models · 248 questions · 4 days

From September 5–8, 2026, we collected **12,504 answers** from 3 AI models. The main experiment used **248 distinct questions**, repeated across 4 days, including commercial service questions and control questions.

We asked both for recommendations and for businesses the models knew about. For example:

“I’m looking for a financial advisor in Denver, Colorado. Who would you recommend?”

“What are some financial advisors in Denver, Colorado that you know of?”

The wording was fixed for each category and question type, with the city substituted consistently. Repeated responses let us examine patterns beyond a single answer.

2

Compare the businesses' websites

3,127

BUSINESS WEBSITES COMPARED

from 5,426 Google Maps listings

We collected **5,426 Google Maps listings** across the 96 local markets, each covering a service in a city. From those, **3,127 businesses** had readable lists of their website pages.

That let us compare websites belonging to businesses AI named with websites belonging to businesses it never named in our sample.

3

Ask the questions buyers consider before hiring

288

ANSWERS TO BUYER QUESTIONS

16 questions · 6 cities · 3 models

On September 10, we asked **16 cost and comparison questions** across **6 cities**, using the same 3 models with web search enabled. That produced **288 additional answers**.

We then classified the cited sources and the types of local-business pages among them.


Investigation	Scale
Main AI experiment	12,504 answers
Website comparison	Page lists from 3,127 business websites
Buyer-question follow-up	288 answers
Total AI answers analyzed	12,792

These are answer counts, not counts of unique questions or individual web searches. Repetition was part of the design. The website investigation was a separate data collection exercise, not additional AI answers.

WHERE WE TESTED

12 metros, chosen deliberately

Local businesses operate in very different markets. We included **4 large, 4 midsize and 4 small metropolitan areas**, or metros, using the study's predefined size groups.



Approximate metro locations. Schematic, not to scale.

LARGE METROS**Chicago**

Illinois

Minneapolis

Minnesota

Atlanta

Georgia

Denver

Colorado

MIDSIZE METROS**Nashville**

Tennessee

Tucson

Arizona

Albuquerque

New Mexico

Rochester

New York

SMALL METROS**Boise**

Idaho

Chattanooga

Tennessee

Peoria

Illinois

Fort Wayne

Indiana

We started with **17 candidate metros**.

Before the main run, we measured how often search results included pages listing businesses, such as “best” or “top” local-service lists. This gave us a consistent, if imperfect, way to compare that type of online coverage across markets.

We chose metros with a range of that measured coverage within each size group. Chicago and Fort Wayne were retained to preserve continuity with earlier pilot work. The final 12-city list was frozen before the main experiment.

This was a deliberately varied sample, not a random sample representing every US business or city.

THE BUSINESS MIX

8 service industries, 4 business groups



PROFESSIONAL SERVICES

Financial advisors
Business attorneys



LOCAL TRADES

HVAC companies
Plumbers



REAL ESTATE SERVICES

Real estate agents
Property management companies



LOCAL PRACTICES

Dentists
Speech therapists

These categories represent services customers hire, including both individual practitioners and companies. Their mix allowed the study to examine different kinds of local service decisions rather than treating one trade as representative of them all.

96**local markets**

8 industries × 12 metros

The broader experiment also included 2 local control categories, auto repair and restaurants, plus 8 national test questions covering consulting, product purchasing and business software.

They helped distinguish answers naming local independent businesses from other kinds of recommendations. The headline findings here focus on the 8 commercial local industries.

The buyer-question follow-up covered Chicago, Atlanta, Albuquerque, Nashville, Fort Wayne and Boise: 2 metros from each size group, across all 8 industries.

AI, EXPLAINED SIMPLY

What “with web search” means

AI can answer using information learned during training. It can also be given access to web search to find information for the question in front of it. We tested both conditions.



WEB SEARCH OFF

11,088

ANSWERS

Existing knowledge

Web search was disabled. The model answered using information learned during training, without looking up websites for that response.

WHAT WE EXAMINED

Which businesses the model could name from its existing knowledge.



WEB SEARCH ON

1,704

ANSWERS

With web search

The model could search online and use the results in its answer, including providing citations.

WHAT WE EXAMINED

Which businesses it named and which pages it cited when search was available.

Why test both? A business appearing in a model’s existing knowledge and a business appearing in its online sources are different routes into an answer. Combining them without labels would hide that distinction.

Condition	Main experiment	Buyer follow-up	Total answers
Existing knowledge	11,088	0	11,088
With web search	1,416	288	1,704
Total	12,504	288	12,792

Search being available does not mean every response used it or included a citation. We count answers produced in that condition, not individual searches performed behind the scenes.

We tested OpenAI GPT-5.6, Anthropic Claude Sonnet 5 and Google Gemini 3.6 Flash through an API, a software connection to the models. The consumer apps manage their own search, prompts and answer presentation, so this is a study of the tested models and settings rather than a direct test of the apps.

FINDING / BUYER QUESTIONS / WEB SEARCH ENABLED

When buyers ask questions, AI cites businesses' answers

70.6%


of cited local-business pages were answer or pricing pages.

Buyers ask what a service costs, how 2 options compare and when a repair makes more sense than a replacement.

In the follow-up experiment, we asked questions like those and examined the sources the models cited.

Of **177 links to local-business pages**, 92 led to answer pages and 33 to pricing or fees pages. **6** led to homepages.

In plain terms: for every 10 local-business pages AI linked to, 7 answered a buyer's question or explained prices and fees. Homepages were almost never the page cited.



The result points to a specific publishing opportunity. Businesses can provide useful explanations of the costs, choices and trade-offs their customers are trying to understand.

For service businesses, that can include explaining what affects a quote, what a fee includes or why 2 solutions suit different situations. A useful cost explanation can go beyond a single advertised price.

Scope: 288 buyer-question answers with web search enabled. Page types were classified from page addresses; these figures do not constitute a manual quality assessment of every page. The finding describes which pages were cited, not the probability that a newly published page will earn a citation.

REAL EXAMPLES, ANONYMIZED

What those answers looked like

The following examples come from saved responses. Business identities, links and exact source-page titles are omitted. Source descriptions are paraphrased; the short AI excerpts retain their original wording.

An air-conditioning cost question

Buyer question: How much does it cost to replace an AC unit in a [large metro]?

Local-business sources cited: A replacement-cost guide and installation-cost articles from HVAC websites.

RECORDED AI EXCERPT

“Final expenses depend on factors such as home square footage, system efficiency (SEER2 rating), labor rates, and whether ductwork modifications or permits are required.”

Content lesson: A page can help a buyer understand what changes the cost, as well as the cost itself.

TO DO

Create a page on your website that breaks down what changes your price, one factor per section, in plain words. When a buyer asks AI what your service costs, give it a specific section to cite.

A property-management fee question

Buyer question: How much do property management companies charge in a [large metro]?

Local-business sources cited: A property-management cost guide and a management-fees page.

RECORDED AI EXCERPT

“Beyond the monthly fee, most also charge a leasing fee for finding tenants”

Content lesson: Explaining the parts of a service helps buyers understand what they are comparing.

TO DO

Break down your service on your website: what the customer gets at each stage and every fee they usually ask about. Give AI a clear breakdown to cite when buyers compare providers.

A water-heater comparison

Buyer question: Tankless or tank water heater: which makes more sense for a home in a [midsize metro]?

Local-business source cited: A plumbing company's tank-versus-tankless comparison page. The answer also cited a manufacturer page and a video.

RECORDED AI EXCERPT

“a traditional tank heater is better if you prefer lower upfront installation costs and simpler maintenance”

Content lesson: Comparison pages can explain which priorities influence a decision.

TO DO

Write a comparison page for the choices your customers weigh, and say plainly who each option suits. That is the page AI can cite when a buyer asks which option is right for them.


These examples demonstrate citation behavior. They do not independently validate the technical advice in an AI answer, or establish that a particular sentence came from one specific source when several were cited.

FINDING / BUYER QUESTIONS / WEB SEARCH ENABLED

36.0%

of cited pages came from local businesses' own websites.

Across the buyer-question follow-up, we recorded **492 page citations**, counting a page once within an answer. Local businesses' own websites supplied **177**, the largest single source category in this classification.



A local business's website can serve 2 audiences at once: a person researching a decision and an AI system assembling an answer to that question.

This is the practical significance of publishing expertise. The information a business provides on its own site can appear among the sources supporting an AI response.

Scope: cost and comparison questions with web search enabled. “Local” used a matching Google Maps website and a listing in the queried city’s state, or a listing without an address. It does not necessarily mean within city boundaries. The same page can count again if cited in another answer. Percentages are rounded.

FINDING / BUSINESS WEBSITE COMPARISON

99 vs 44

99 content pages versus 44.

We compared the lists of website pages for **758 businesses named by AI** and **2,369 businesses never named** in the study sample.

The typical business AI named had **99 content pages**, compared with **44** for a business it never named. It also had **5 pages addressing questions or comparisons**, versus **1**. These figures use the middle business in each group, known as the median.

Median content pages



Median question-style pages

Named by AI 5

Never named in the sample 1

“Median” means the middle business in each group. It prevents a few very large websites from dominating the comparison.

Question-style pages were identified from addresses suggesting a buyer question or comparison. Property listing feeds were excluded from content counts, so a large inventory of homes was not treated as a large library of advice.

The named businesses tended to have built more on their websites. But website size can also accompany business size, age and prominence. That is why we went beyond this first comparison.

Scope: businesses with readable sitemaps, combining both AI conditions. “Never named” means absent from this experiment, not invisible to every AI system. This is an observed association, not evidence that doubling page count doubles recommendations.

FINDING / SAME SERVICE, CITY AND SIMILAR REVIEWS


The content difference remained with similar review counts

59.6%

With similar review counts, the business AI named had more website pages in roughly 6 of 10 comparisons.

Businesses with more established websites may also have more reviews. To make the comparison closer, we paired named and never-named businesses in the same category and city, keeping **348 pairs** whose review counts were within 25% of one another under the study's matching rule.

Within those pairs:



Excluding the 4 pairs with equal page counts, the business AI named had more pages in **59.6%** of comparisons. The difference was statistically significant ($p = 0.0004$).

This is stronger evidence than the unmatched comparison: the content association remained when businesses had similar review counts and competed in the same market.

The clearest result came from businesses named using **existing AI knowledge**. The separate comparison for businesses named only with web search did not establish a statistically significant content advantage.

For business owners, the finding supports paying attention to the depth of the website alongside reviews. It does not identify a magic page count or prove that adding pages will cause AI to name a business.

Matching reduces one obvious alternative explanation. It does not account for every difference between businesses, such as age, brand prominence or third-party coverage.

FINDING / EXISTING KNOWLEDGE / WEB SEARCH OFF

7 of 96

The business with the most reviews was also AI's most-mentioned business in only 7 of 96 markets.

In the test with web search turned off, we checked whether the business with the most Google reviews was also the business AI named most often. They matched in **only 7 of the 96 local markets**.

We also examined review counts and mention frequency across **2,533 verified businesses with Maps data**, accounting for differences between categories. Review count explained **4.5% of variation in how often businesses were named**.

That does not mean reviews account for 4.5% of an AI ranking formula. It measures the strength of the observed relationship in this analysis.

The useful finding is that collecting the most reviews did not reliably make a business the model's most-named choice. Review totals alone leave much of the pattern unexplained.

Reviews still matter in parts of the data, particularly when comparing whether businesses were named at all. This result concerns mention frequency among verified named businesses, with web search disabled. It is not a reason to stop collecting customer feedback.

But **a review strategy alone is an incomplete explanation of AI visibility**. The website comparisons provide another part of the picture, without establishing that content accounts for all the unexplained variation.

FINDING / BUYER QUESTIONS / WEB SEARCH ENABLED

55%

of the local businesses AI linked to weren't in the first 60 Google Maps results we checked.

When AI answered buyers' cost and comparison questions, it linked to websites belonging to 129 local businesses. **71 of them (55%) did not appear in the first 60 Google Maps results** we collected for that service and city.

We checked the first 3 Maps pages for the corresponding service and city.



This demonstrates that an AI citation can reach beyond the businesses leading a particular Maps query. A business outside those results could still have a page cited when the question concerned costs or choices.

It does not establish that Maps visibility is irrelevant, or that publishing content will bypass every other factor. Maps results depend on the query, location and collection time; the citation check also allowed businesses elsewhere in the same state.

The opportunity is observable: useful business pages appeared as AI sources beyond the leading Maps listings we collected.

READ THE RESULT CORRECTLY

Being named and being cited are different

There are 2 outcomes in this report that a business should keep separate.

Named in the answer

WHAT THE BUYER SEES

A business appears in the response text, such as a provider offered for consideration.

WHAT WE MEASURED

Whether a local independent business was named, and how often.

Cited as a source

WHAT THE BUYER SEES

A business page appears as a supporting link.

WHAT WE MEASURED

Which pages and businesses supplied citations.

WHEN WE ASKED WHO TO HIRE

57.9%

of own-business-site citations went to homepages

WHEN WE ASKED COST AND COMPARISON QUESTIONS

70.6%

of local-business page citations went to answer or pricing pages

Those are different question sets and citation samples. Their contrast shows why a website needs to be considered in the context of the buyer's question.

A link to a business's website is not automatically a recommendation to hire that business.

WHAT TO PUBLISH

What this means for your website

The clearest publishing opportunity in this study is to answer the questions people ask while deciding what to do.

Start with the questions your customers already bring to conversations. What affects the price? What does the service include? Which option fits their situation? When should they repair, replace, hire or wait?

Use the business's actual experience to make those explanations specific. Include the distinctions customers regularly miss. Where advice depends on location, circumstances or professional assessment, explain that context.

For businesses that do not publish fixed prices, useful cost content can explain the pricing model and quote variables. The study observed citations to both answer pages and pricing pages; it did not test whether this particular alternative earns the same results.

Keep a clear homepage as well. The recommendation questions and research questions drew on different parts of business websites.


These are practical applications of the findings, not a tested publishing formula. The evidence supports creating substantive buyer-facing content; it does not set a required article count, posting frequency or guaranteed outcome.

ANSWER CONTENT ENGINE

We're the best AI visibility data nerds you'll find

We analyzed 12,792 AI answers to write this report. Our mission is to get your business cited and named by AI as the answer when your buyers ask.

7 agents do the work. They find the questions your buyers ask, rank them by buying intent, write each answer with your proof, create the images and video, publish the moment you approve, improve what's already working and track where you show up against competitors.



Done for you. You just approve.

About this report: Liron Segev, who built Answer Content Engine, conducted the study. This was research into AI recommendations and citations, not a trial measuring Answer Content Engine customer results.

CLIENT RESULTS

We get results

This is what results look like for some of our clients:

64%

of AI category mentions for a solo real estate agent, up from zero.

71%

of category mentions for a sports consulting business, nearly 3x its nearest competitor.

#1

most-mentioned storage facility in its category: 38 AI mentions, 32 first-page Google spots.

Anonymized client results from live Answer Content Engine dashboards, April to June 2026. They are not results from this study.



The Authority page in your portal tracks AI mentions, Google top-10 placements and your share of category mentions every week. Competitor names hidden.

NEXT STEP

Your customers are asking AI questions. You need to be the answer.

✔ You own the system

✔ Not a SaaS tool

✔ Built by data nerds

Book a call to see how it works for your business.

Book a call here →

AnswerContentEngine.com



★★★★★

"One of my newest clients found me because Claude cited my content... It's insane to see the results coming in so much faster than I ever expected."

Andrew Murdoch, Founder, YT Era · Google Review

Don't wait until your business is completely invisible in AI. Is it not worth a 15-min chat?

METHOD AND INTERPRETATION

Method and interpretation

Study scope

The main experiment collected 12,504 answers across 4 days, September 5–8, 2026. The website investigation and 288-answer buyer-question follow-up were completed on September 10. Test runs were excluded from the reported answer totals.

The 3 models were Anthropic Claude Sonnet 5, OpenAI GPT-5.6 Sol and Google Gemini 3.6 Flash, accessed through OpenRouter. We used each provider's default settings for how the model reasons and generates answers. The consumer apps were not directly tested.

Geographic selection

17 candidates were assigned to the study's large, midsize and small metro groups. Before selection, we measured titles suggesting business lists among up to 100 organic search results per service-and-city query. We put those measurements on a common scale and combined them across 7 commercial categories. Business attorneys remained a full study category but were excluded from the selection measure because search results were incomplete or inconsistent. We selected a range of scores within each size group; Chicago and Fort Wayne were retained from the pilots. These are study groupings, not a claim to official population thresholds or national representativeness.

Repetition and controls

The main experiment used fixed recommendation and existing-knowledge questions, with consistent city substitutions, plus control questions. Repeated answers are repeated observations of the same questions, not independent samples of US markets. The 12,792 total includes controls and both search conditions; each finding identifies the sample it uses.

Business names and verification

Distinct business names from the existing-knowledge condition were checked against Maps, with further web checks for unmatched names. Identity matching accounted for naming variations. The web-search recommendation names did not receive the same full verification; some website matches used names, which introduces uncertainty.

Website comparison

The sample began with Maps pages 1–3 across 96 commercial markets. Website inventories were readable for 3,127 businesses. Chains, directory or brokerage sites, and ambiguous identity matches were removed under the analysis rules. Content counts excluded property listing feeds. Page types were inferred from addresses. The main close-review comparison retained 348 pairs; equal content-page counts were excluded from the percentage calculation.

Citation follow-up

16 cost and choice questions, 2 per commercial industry, were asked in 6 cities to 3 models: $16 \times 6 \times 3 = 288$ answers. The analysis recorded 492 page citations, including 177 classified as local-business pages. Repeated links within an answer were counted once; a page cited in different answers could count again. Local classification used Maps website matching and state-level geography, with an allowance for listings without addresses. Government, nonprofit and media listings were excluded from the local-business group.

Quality checks

The main classifier was checked against 100 held-out answers labeled by hand. Agreement was 88.0% across answer categories and **96.0%** on whether an answer named a local independent business. The final datasets reported no errors or unclassified answers. These checks assess processing reliability; they do not certify every statement generated by the models.

What the evidence can establish

The website findings are observational. Matching reviews improves comparability but cannot isolate the causal effect of content. The content advantage was clearest for existing-knowledge naming; the search-only comparison was inconclusive. The buyer-question follow-up showed the types of pages cited, but its 43 reviews-matched business pairs were insufficient to establish that more content causes more citations.

This report focuses on content, reviews and citation opportunities within a broader experiment. Findings outside that focus and detailed analytical notes are retained in the full study record.