AndrewRyanMarketing

THE ANSWER GAP INDEX™ · EDITION 2 · OCTOBER 2026

Who AI recommends is not who Google ranks.

Edition 2 · 720 AI answers · 4 engines · 12 business types · 3 cities · 5 runs each

We asked Google's AI Overview, Google's AI Mode, ChatGPT and Perplexity to name the best local business: the same question, five times each, for 12 kinds of business in Denver, Boise and Chattanooga. Then we checked every business they named against Google's local map pack and organic top 10 for the same search. Edition 2 repeats Edition 1 one month later on the same 720 questions and tests each of the six measures for change. The Answer Gap Index™ reports six measures, not a single score.

GHOST WINNERS™

58.3%

of the 640 businesses AI recommended in October were in **neither** Google's map pack nor its organic top 10, in any of five captures of the same search. September: 60.5% of 673. No detectable change.

What changed since Edition 1

Measure	Engine	September	October	Change (95% CI)	Result
Ghost Winner Rate	All engines	60.5%	58.3%	−2.2 (−6.9, +2.7)	**held**
Ghost Top-Pick Rate	ChatGPT	36.0%	40.5%	+4.4 (−8.8, +18.8)	**held**
AI Overview Trigger Rate	Google	36.1%	27.2%	−8.9 (−20.0, +3.3)	**held**
Top-Answer Lock Rate	ChatGPT	27.8%	38.9%	+11.1 (−11.1, +33.3)	**held**
Cross-Engine Agreement, #1	3 engines	6.2%	3.6%	−2.7 (−9.2, +2.3)	**held**
Evidence Ownership Rate	ChatGPT	6.9%	**26.9%**	**+19.9 (+13.4, +27.1)**	**rose**

A change counts only if the 95% confidence interval of the October-minus-September difference, from a cluster bootstrap over the 36 city-and-business-type cells (2,000 draws, both editions resampled on the same cells), excludes zero. Otherwise the measure is reported as held. The rule was written on 28 September 2026, before Edition 2 was collected. Points are percentage points.

Persistent Ghost Winners™ (new in Edition 2): 71% of October's Ghost Winners (147 of the 207 that appeared in both editions) were already Ghost Winners in September. Being recommended by AI while absent from the rankings is mostly a standing condition, not a one-day fluke.

The six measures, by engine

Engine	Businesses named	Ghost Winners	#1 pick is a Ghost Winner	Same #1 every run	Answer cites own site
ChatGPT	248 Sept 310	53.2% Sept 58.1%	40.5% Sept 36%	39% Sept 28%	**27%** Sept 7%
Google AI Mode	357 Sept 372	56.9% Sept 53.2%	37.8% Sept 28%	6% Sept 3%	12% Sept 21%
Perplexity (API)	210 Sept 244	56.2% Sept 49.6%	42.0% Sept 35%	72% Sept 75%	62% Sept 77%
Google AI Overview	152 Sept 148	30.3% Sept 23.0%	23.7% Sept 8.5%	n/a	29% Sept 54%
All engines	640 Sept 673	58.3% Sept 60.5%			35% Sept 37%

October 2026 values, with September 2026 beneath. Unique businesses per city and engine; businesses we could not resolve to a website (49 in October, 60 in September) are excluded, which keeps the Ghost Winner Rate conservative. “Same #1 every run” = share of 36 city-and-business pairs where the first business named was identical in every run; the AI Overview never appeared in all five runs for any pair, so it has no lock rate. “Answer cites own site” is counted per recommendation, across all runs. Perplexity, Google AI Overview and Google AI Mode exposed fewer citations per answer in October than in September (Perplexity about half as many, AI Mode about a quarter), which lowers that per-answer measure mechanically; their lower October values are reported but not counted as changes. Only the ChatGPT rise passed the change test.

The six findings

01 **Ranking helps. It doesn’t decide who gets named.**
53.2% of the businesses ChatGPT recommended held neither a map-pack nor a top-10 spot (September: 58.1%). Across all four engines the Ghost Winner Rate held at about 60%: 58.3% in October against 60.5% in September, across home services (60.9%), professional services (52.7%) and B2B (61.6%).

02 **The #1 pick is often a Ghost Winner.**
The first business named held no map-pack or top-10 spot in 40% of ChatGPT answers, 42% of Perplexity’s and 38% of AI Mode’s; none of these moved more than chance allows. The AI Overview’s first pick was in the map pack about half the time (73% in September).

03 **Google’s AI Overview mostly sits these searches out.**
It appeared on 27% of “best [business] in [city]” searches (49 of 180), against 36% in September: a drop inside the noise. Google’s AI Mode answered 97% of the time, and 56.9% of the businesses it named were Ghost Winners.

04 **The top spot rarely holds.**
Asked the same question five times, ChatGPT named the same #1 every time in 39% of pairs (28% in September); AI Mode in 6%. Perplexity’s API held 72%, at default settings.

05 **The engines don’t agree with each other.**
ChatGPT, Perplexity and AI Mode put the same business first in 3.6% of 167 runs (6.2% in September), and named no business in common at all in 63%.

06 **ChatGPT now cites the business’s own website for one pick in four.**
27% of ChatGPT’s recommendations came with a citation to the business’s own website, up from 7% in September: the one measure that moved more than chance allows (95% CI +13 to +27 points), and it moved the same way in every business group. Most of the proof still lived on directories, lists, reviews and forums.

How we measured it

Same prompt on every engine: “best [business type] in [city, state].” Five runs per city, business type and engine on 27 October 2026 (720 answers), collected through the DataForSEO API; US English, desktop. Identical in every setting to Edition 1 (28 September 2026), under a protocol written before Edition 2 was collected. Google results were located to each city; ChatGPT and Perplexity to the US with the city in the question; Perplexity via its API (sonar, web search on). Named businesses were matched to their own websites and checked against five captures of the map pack and organic top 10. The 49 we could not resolve (7%) are excluded, which keeps the Ghost Winner Rate conservative. Edition 2 also captured Google from four more points about six miles north, south, east and west of each city centre on the same day: from all five points, 50.0% of AI-recommended businesses were still outside Google’s top-3 map box and top 10. One day per edition, three cities, desktop only: two editions a month apart are the first two points of a benchmark, not a trend. None of the businesses in the answers are Andrew Ryan Marketing clients.

Where does your business stand?

Uncover Your Revenue Gap: free, three minutes, no email.
andrewryanmarketing.com/6-gap-revenue-map

Full study, charts and city-level data for both editions: **andrewryanmarketing.com/answer-gap-index**
© 2026 Andrew Ryan Marketing. Licensed CC BY 4.0: free to share and adapt with credit and a link to the study.