Who Do AI Answers Cite Most?

49 brands, 235 answers, 91 mentions - 12-19 November 2026

We asked Gemini and ChatGPT buyer questions about 49 real brands and saved every answer in full: 235 answered questions, 91 mentions, 7 calls that never returned an answer. Caraway and Our Place lead the count with 8 mentions each. Eleven brands were mentioned every time they were asked. Six brands were never mentioned once. And there is one question this report cannot answer, stated plainly in section 2: which kind of page gets cited most - because the cited sources were never saved, and this report does not guess at what it did not record.

Online version: https://auditlab.site/blog/who-ai-answers-cite-most

Cite as: GeoSeen, "Who Do AI Answers Cite Most?", 12-19 November 2026. Quote any number with a link back to the page above and the date alongside the number.

The headline numbers

One observation is one brand x one question x one engine with an answer. A call that failed (HTTP 429/503, or no answer) is missing - never counted as a zero, never in the denominator. The same question re-asked on another day is a separate observation. Snapshot: 12-19 November 2026, leaderboard data updated 19 November 2026. Gemini contributed 213 answered questions and 83 mentions; ChatGPT contributed 22 answered questions and 8 mentions. Claude contributed nothing: it was blocked in every batch, so nothing here is a three-engine claim.

Measure	Count
Brands tested	49
Answered questions (observations)	235
Mentions	91
Missing (failed calls, excluded)	7
Overall mention rate	91 of 235 - 38.7%
Brands mentioned at least once	43 of 49
Brands never mentioned	6 of 49

Read the 38.7% as what it is: across all the buyer questions we asked, a brand was named in roughly two of every five answers. It is not a market share and not a Google ranking. It is a count of saved answers.

1. Who gets cited most - by count

This report ranks by mention count. The leaderboard ranks by mention rate, and the free summary follows the leaderboard - same data, a different cut. Count and rate tell different stories: a brand asked 21 times can lead the count while a brand asked twice can lead the rate.

#	Brand	Mentions	Answered	Mention rate
1	Caraway	8	21	38.1%
1	Our Place	8	21	38.1%
3	Brooklinen	7	21	33.3%
4	Ghost	4	4	100%
4	ME+EM	4	4	100%
4	Whistles	4	4	100%

7	Outer	4	21	19.0%
8	Boohoo	3	4	75.0%
8	Notion	3	4	75.0%
10	Twelve brands tied at 2 mentions (see no	2	-	-

Tied at 2 mentions: Balsamiq, Calendly, Clockify, Loom, Medik8, Neal's Yard Remedies, Tally and Todoist (2 of 2 each); Allbirds, Brightland, Figma and Warby Parker (2 of 4 each).

Three things to read out of this table, and one warning. First, the top three by count - Caraway, Our Place and Brooklinen - hold 23 of the 91 mentions between them, 25.3% of everything cited. Second, the perfect records sit below them: 11 brands were mentioned every time they were asked - Ghost, ME+EM and Whistles on 4 answers each, and Balsamiq, Calendly, Clockify, Loom, Medik8, Neal's Yard Remedies, Tally and Todoist on 2 answers each. A 100% rate on 2 answers is a small-sample fact, not a crown; say the denominator with the rate or the rate means nothing. Third, the warning: counts follow question volume. Five brands - Caraway, Our Place, Brooklinen, Outer and Faherty - were scored on the weekly leaderboard batch of 20 mixed questions each, so 107 of the 235 answered questions (45.5%) belong to just those five. A brand asked 21 times has 21 chances to be mentioned; a brand asked twice has two.

Boohoo's row carries a footnote the table cannot show: 3 mentions in 4 answered questions, with 2 further calls missing (failed, excluded - not counted as zeros). Its 75.0% is 3 of 4, not 3 of 6.

2. What kind of content gets cited most

NOT VERIFIED - the data does not support an answer. The 235 answers recorded whether each brand was mentioned and, in about half the cases, where it sat in the answer's list. They did not record the source URLs the engines cited. Across the 23 source files behind the leaderboard tally there are zero cited-source entries - the URLs that do appear in those files are site checks (llms.txt, robots.txt, sitemaps) and chat links, not the sources behind the answers. So "which page type gets cited most" - brand home page, product page, review roundup, marketplace listing, Wikipedia - cannot be ranked from this data, and this report does not rank it. Any report that did would be inventing the ranking.

Two things the data does support, and one change for next time. First, list position: of the 91 mentions, 50 have a recorded position inside the answer's own list - 22 at #1, 16 at #2, and 12 at #3 or lower. The other 41 mentions have no recorded position in the source files, so the position split describes the 50 we can place, not all 91. A mention is usually a list item, and when we can place it, it is usually near the top - but 41 unplaced mentions is too many to call that a finding about all mentions. Second, the llms.txt check: five brands have a verified llms.txt status from the same days, and the pattern runs against the file. Everlane (llms.txt live, 200, 5,793 bytes) scored 0 of 4. Faherty (live, 200, 5,730 bytes) scored 0 of 23. Raycast (live, 200, 27,002 bytes) scored 0 of 4. Figma (no llms.txt, 404) scored 2 of 4. Caraway (no llms.txt, 404) scored 8 of 21. Five brands is a small sample and this is an observation, not a conclusion - but in this sample, having a real llms.txt predicted nothing about being mentioned, in either direction. Third, the fix: from the next batch on, every saved answer keeps its cited-source list (Gemini grounding sources, ChatGPT sources) next to the mention count, so the next edition of this report can rank page types from recorded sources instead of declaring the question unanswerable again.

3. The six never mentioned

Six of the 49 brands were not named in any answer we recorded. The denominators differ because the brands were asked different numbers of questions - read the fraction, not just the zero.

Brand	Mentions	Answered	Missing	Engines
Albaray	0	6	0	Gemini
Everlane	0	4	0	Gemini, ChatGPT

Faherty	0	23	1	Gemini, ChatGPT
Finery London	0	4	0	Gemini
Raycast	0	4	0	Gemini, ChatGPT
Tropic Skincare	0	2	0	Gemini

Faherty's 0 of 23 is the most-tested zero in the set - 21 of those 23 are the weekly leaderboard batch, 20 mixed questions that include ones outside clothing, plus 2 targeted questions. Everlane's 0 of 4 and Raycast's 0 of 4 are two-engine zeros (Gemini and ChatGPT), not three-engine ones: Claude was not tested for any brand here. Tropic Skincare's 0 of 2 is the thinnest zero - two questions, one engine, one day. And a zero is a snapshot of 12-19 November, not a verdict on the brand: in our own earlier test a zero flipped overnight when the same question was re-asked the next day. What a 0 of 23 does tell you, narrowly, is that across 23 saved answers on two engines, buyers asking those questions were shown full shelves with Faherty missing from all of them.

What to do with this

Three checks, each tied to a number above.

- [C-01] Count with the denominator attached: a mention count without its answered count is a vanity number. Caraway's 8 sits on 21 questions; Ghost's 4 sits on 4. Quote both numbers together or the comparison is meaningless.

- [C-02] Save the sources, not just the mentions: if you run mention checks, keep the cited-source URL list with every answer. Without it you can say who was mentioned and never say what kind of page earned the citation - the exact gap in section 2 of this report.

- [C-03] Re-ask a zero before acting on it: a zero from one day's batch is a prompt to re-ask, not a fact to print. Faherty's 0 of 23 survives re-asking across 21 questions; Tropic Skincare's 0 of 2 has only been asked twice.

Method note - read this before quoting the numbers

Sample: 49 brands, 12-19 November 2026; 23 source files; 235 answered questions, 91 mentions, 7 missing. One observation is one brand x one question x one engine with an answer. Missing (HTTP 429/503 or no answer) is excluded from the denominator and never counted as 0. The same question re-asked on another day is a separate observation. Engines: Gemini (213 answered, 83 mentions, free API key) and ChatGPT (22 answered, 8 mentions, free account, already signed in). Claude: no observations - blocked in every batch, nothing claimed about it. Mention means the brand name occurs in the full saved answer text; list position means the position inside the answer's own list or section, where recorded (50 of 91 mentions have one). File facts (llms.txt status and byte size) are same-day fetches with HTTP status recorded. GeoSeen itself is listed separately on the leaderboard, not ranked: one of its three questions names GeoSeen in the question itself. Every count in this report comes only from completed answers with full text saved; no failed call was ever counted as a zero. Snapshot, not a trend: re-ask before acting on any single zero, including the ones in this report. Live totals match the public API (/api/public/leaderboard, updated 2026-11-19) exactly: 49 brands, 235 answered, 91 mentions, 7 missing.

GeoSeen - auditlab.site - November 2026. Free to quote with a link and the date.