

Can a Vendor's Own "Best Tools" List Earn AI Citations If It Discloses Who Wrote It?

The short answer is yes. A disclosed vendor comparison can earn AI citations, because the engines are barely scoring authorship at all. In the largest count anyone has published, between a third and a half of the "best tools" lists cited in AI answers were written by a vendor selling in that category. The traits that measurably raise a page's citation odds are a comparison table and a stated selection method — not who published the page.

What got counted

In August 2026, the agencies Minuttia and Convex Works asked 18,084 software-buying questions across six platforms — Google Search, AI Overviews, ChatGPT, Gemini, Claude, and Perplexity. They then read all 311,766 pages those platforms cited and labelled each one by who published it and what kind of page it is. It is the first count of vendor-written lists in AI answers that works from the cited pages themselves rather than from opinion about them.

The study's own caveats deserve the same prominence as its findings. Every share it reports is a floor, because pages its team could not read were counted as not self-promotional. And the results are a snapshot, not a standing: when the researchers re-asked Perplexity 2,000 of the questions a month later, only 28% of the pages in the new answers had been there in August. Citations churn. Anyone measuring them from a single screenshot is reading weather and calling it climate.

How much of the answer is a vendor grading its own homework

For "best alternatives to X" questions, 34% to 47% of the listicles the platforms cited were a vendor's own list, depending on the platform — 47% on Google and Gemini, 34% on ChatGPT. For "how do X and Y compare" questions, the vendor-written share rises to between 56% and 65%. Across every category and question type, 43% of cited listicles are self-promotional.

The spread by category is wider than the average suggests, running from 12% in audio production to 76% in AI security operations. AI visibility — the category this site sells in — sits at 65%, one of the most self-promotional categories measured. If you buy or sell in a software category, the honest working assumption is that a large share of the "best tools" content the machines quote was written by someone ranked in it.

The lists themselves behave the way you would expect. Of the vendor-written lists the platforms cited, 93% highlight the vendor's own product, 87% rank it first, and 90% criticise competitors. Only 16% clearly disclose that the vendor wrote the page. The typical cited list, in other words, is an undisclosed advertisement that puts its author first — and the engines cite it anyway.

Each engine shops differently

ChatGPT searches the way a buyer going straight to the source would. In 71% of the searches it ran before answering, it named exactly one product, and 92% of its site-limited searches pointed at a vendor's own website. On alternatives questions, vendors' own sites made up 42% of ChatGPT's citations, against 1% to 6% on the other platforms. For ChatGPT, the vendor's own pages are not a contamination of the research — they are a primary input to it.

Perplexity works more like an editor with a pile of candidates. A typical answer finds about 20 pages and cites 9 of them, and which pages survive depends first on where they sat in its search results. Here is the finding that reframes the whole question: of the pages Perplexity found, vendor-written listicles were cited 58% of the time — the highest citation rate of any page type it found, ahead of independent listicles at 44% and review sites at 36%.

There is no single shelf to win, either. Only 21% to 28% of the citations from AI Overviews, Gemini, Perplexity, and Claude also appear in Google's top 10 results for the same question. Each engine assembles its own evidence, which means a page earns its place engine by engine, on traits each engine can see.

The traits that actually move a citation

Perplexity is the one platform where the study could watch the selection happen — every page found, then which ones got cited. Holding search position equal, here is how often a found listicle was cited with and without each trait:

Trait	Cited with it	Cited without it
Has a comparison table	70%	60%
Says how the tools were picked	69%	61%
Longer than the typical list	68%	62%

Trait	Cited with it	Cited without it
Year in the title	68%	63%
Written by a vendor in the category	64%	68%
Promotional tone	64%	69%
Ranks its own product first	64%	65%

Read the table from the bottom up and the story tells itself. Ranking your own product first — the move that looks most like cheating — costs about one point. A promotional tone costs about five. Being a vendor at all costs about four. Now read the top: a comparison table is worth ten points, and stating how the tools were selected is worth eight. The penalties for self-interest are small and the rewards for being genuinely useful are larger, on the same scale, in the same measurement.

One honesty note about this article's own title: the study never isolated disclosure itself as a citation trait, so no one can honestly claim a measured disclosure bonus. What the study did measure is disclosure's closest cousin — a page that says how its tools were picked gets cited more, not less. Nothing in the data shows an engine detecting a disclosure line and punishing it, and the 84% of vendor lists that hide their authorship are not winning because they hide. They are winning in spite of it, on tables, method, length, and position.

So why disclose, if the machines don't make you

Because the citation is only half the transaction. The other half is the human who clicks through from the answer and decides whether to trust the page in front of them. A list that hides its author converts every citation it earns into a small liability: the reader who discovers the trick discounts not just the list but everything else the vendor says. Disclosure costs no measurable citation share, and it is the only version of the page that survives both readers — the machine counting traits and the person counting on you.

This is also a game a small operator can actually win. Vendor-written lists that earn citations come from smaller sites than the independent lists they compete with — a median domain rank of 49 against 63 — and the field is fragmented: the 20 most-cited vendors together hold only 6% of the citations. There is no incumbent wall here. There is a craft bar, and it is published in the table above.

The rules we publish by

- **Disclose at the top, not in the footer.** If a page compares a category we sell in, the first screen says so. A disclosure a reader has to hunt for is a hiding place with better manners.
- **State the selection method.** How the tools were picked is a measured citation trait, not just an ethical nicety. Write the method before the rankings, and let the rankings follow from it.
- **Put the comparison in a real table.** Rows a machine can parse and a person can scan. The table is worth more citation share than any headline about yourself.
- **Earn first place or don't take it.** Ranking yourself first costs almost nothing in citations and almost everything in credibility when the method on the same page doesn't support it.
- **Date it and keep it current.** A year in the title is itself a small citation trait, and a stale list quietly stops being evidence. Lists are maintained assets or they are litter.
- **Publish the data behind the claims.** Minuttia released the data behind every chart in its study so competitors could check the work. Checkable beats confident, with machines and with people.
- **Measure citations as a 30-day rate.** A quarter of the cited pages turn over in a month. A rate over time is the only reading that distinguishes a page that earns its place from a page that got lucky once.

The engines have already decided the underlying question: vendor lists are sources now, cited at rates no one can opt out of by refusing to write them. The live question is narrower and more useful — whether your list reads as evidence or as advertising when it gets there. The disclosed list, with its method stated and its comparison in a table, is the one built to be both.

Study: Minuttia × Convex Works, "Self-Promotional Content Data Study" (August 2026 data), self-promotional-content-data-study.minuttia.com — full data published with the study. Coverage: Digital Market Reports, October 22, 2026.