CATEGORY GRAVITY · SYNTHESIS 006

A short field guide for checking the sources behind an AI answer, then choosing one evidence gap your company can close.

.

WHY AI RECOMMENDS YOUR COMPETITOR

Companion to The AI Trust Gap, Petrichor's August 20, 2026 probe of 15 B2B buyer questions across five AI engines. This guide is a way to inspect your own answer space. It is not a promise that one page or schema change will win a recommendation.

Start with the buyer's question

Your prospect may ask an AI assistant for a category recommendation before visiting your site. The answer can name a competitor, cite an industry source, or give a generic framework that never reaches your company. Those are different problems. Record the answer and its sources before deciding what to publish.

In our dated probe, ChatGPT and Claude produced 345 citations across 189 domains. McKinsey appeared 19 times, Andreessen Horowitz 14, and Y Combinator 12. Petrichor appeared once. The sample is a fixed set of positioning and distribution questions, not a forecast for your market. Perplexity and Gemini rate-limited during capture, so we did not treat their missing results as zero. Read the method and dataset.

The three-question citation check

Choose **three real questions** a buyer asks when the problem is active: one about the problem, one about options, and one about a decision. Use the buyer's words. If you only have guessed questions, mark them as hypotheses until you confirm them in interviews or sales calls.

1. **Problem:** What is causing [specific costly symptom] in [our situation]?
2. **Options:** What are the ways to solve [specific problem], and when does each make sense?
3. **Decision:** Which provider or approach should [buyer role] consider for [use case], and what evidence supports the choice?

Skip to content

Run the same wording in the engines your buyers use. Record the date, model, answer, named companies, cited URLs, and whether each URL actually supports the claim in the answer. Repeat each question at least three times. AI answers vary between runs; a single response is an observation, not a baseline.

Buyer question and engine	Named company or approach	Cited URL and what it supports	Gap we can verify
1. Problem:			
2. Options:			
3. Decision:			

Print or save this page as a PDF to use the blank table in a working session. For a repeatable audit, add a row for every run and keep the exact prompt and response.

Read the result before writing

If a competitor is named, inspect the evidence behind that mention. A cited page may contain original measurement, a clear product description, or only a recognizable name. If the answer cites a third party instead, ask whether that source has evidence your company can contribute. If the answer cites no source, do not count the mention as a verified recommendation.

Make a small evidence inventory:

Check	What to look for	Your finding
Answer presence	Does the company appear in the answer, a citation, both, or neither?	
Source quality	Does the linked page directly support the claim made?	
Original proof	Which claim could we substantiate with our own data or customer evidence?	
Entity clarity	Are company name, offer, people, and canonical URLs consistent across owned pages?	
Third-party corroboration	Which credible external source independently describes or cites our work?	

Choose one move

If the answer lacks useful evidence, publish the piece you can uniquely substantiate: a dated measurement, a documented method, or a case with a clear scope and permission to share it. Show the underlying sample and limits.

If the evidence exists but is hard to find, give it a stable public URL, a precise title, a named author or organization, and links from relevant pages. Keep the same company and offer description across your site and public profiles.

If a trusted source already discusses the problem, contribute a verifiable finding it can cite. A third-party citation is earned through useful evidence and editorial judgment, not through a requested mention.

Skip to content

Then rerun the same prompts on a dated cadence. Compare answers and cited sources by engine. Note changes in model behavior and prompt wording before calling a movement a win.

Decision record

- **Buyer and problem:** _____
- **Question set confirmed by:** _____
- **Baseline date, engines, and runs:** _____
- **Most consequential missing or weak source:** _____
- **Evidence we can publish or improve:** _____
- **Owner and due date:** _____
- **Date to rerun the same questions:** _____

The research report explains what our fixed probe measured. If you want the same check run against your own buyer questions and competitive set, talk to Petrichor.