Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

HOW AI WORKS

How Reasoning Models Decide Who to Cite

By Scott Tischler · October 14, 2026 · 9 min read

Reasoning models — the ones that "think" before they answer — don't just retrieve a few sources and paraphrase the top one. They draft an internal line of reasoning, test it against what they found, discard sources that don't hold up, and only then write the answer you see. The practical upshot: the citation no longer goes to whoever ranked first. It goes to whoever survives the model's own scrutiny. If you want to be cited by a reasoning model, you have to be the source that holds up when the model checks its work.

Over the past year, the frontier models from every major lab added some version of this deliberation step. ChatGPT, Gemini, Claude, and Perplexity's reasoning modes all do a version of the same thing: they spend extra compute working through a problem before committing to an answer. This is a different thing from "deep research" mode, which goes out and runs many searches over several minutes. Reasoning can happen on a single, ordinary query. And it quietly reshapes who gets named.

I run AIrecommend.ai, and this is one of the shifts I watch most closely, because it's invisible. You can't see the thinking step. You only see the answer. But the answer is downstream of a process that is far more skeptical than the keyword-matching retrieval most people still optimize for.

What a reasoning model actually does before it answers

Strip away the branding and the mechanics are roughly consistent across models. The model takes your prompt, often rewrites it into one or more cleaner internal questions, retrieves candidate material, and then — here's the new part — generates an intermediate chain of reasoning that it does not show you by default. In that hidden step it's effectively asking itself: Does this source actually answer the question? Do these sources agree? Is anything here contradicted? What's the most defensible answer I can give?

A few behaviors fall out of that process that matter for visibility.

The model weighs consistency. When several sources say the same specific thing, that claim gets more weight in the reasoning trace. A single page asserting something no one else corroborates tends

to get flagged as weak and dropped, even if it ranked well. This is the corroboration effect, and reasoning amplifies it, because the model is explicitly checking whether its sources line up.

The model resolves contradictions instead of ignoring them. Older retrieval would grab the top chunk and run with it. A reasoning model that finds two sources disagreeing will often try to adjudicate — and the source that loses that adjudication doesn't get cited. Being present in the index isn't enough. You have to be the version the model decides is right.

The model prefers sources that let it show its work. If a page states a claim and also gives the reasoning or evidence behind it, the model can fold that reasoning into its own trace. A page that just asserts a conclusion gives the model nothing to reason with. My read — and this is a read, not a guarantee — is that explained claims travel further through the thinking step than bare assertions do.

Why reasoning changes the citation game

The old mental model for AI visibility was basically search with a paraphrase on top: rank high, get summarized, maybe get a link. Reasoning breaks that model in a specific way. It inserts a judgment step between retrieval and output, and that step is adversarial toward your content. It is actively looking for reasons to not trust a source.

That sounds hostile, but it's actually clarifying, because it tells you exactly what to optimize for. You're no longer optimizing to be *found*. You're optimizing to be *believed* under scrutiny. Those are different jobs.

Being found rewards breadth, keywords, and freshness. Being believed rewards specificity, internal consistency, corroboration from other sources, and claims that are falsifiable enough to check. A reasoning model treats a vague, hedge-everything page the way a careful reader does — with suspicion. It treats a page that commits to a clear, verifiable claim, and shows why, as something it can lean on.

There's a second-order effect worth naming. Because the reasoning step can discard sources mid-process, your competitors are now in the room with you even when the query isn't comparative. If the model retrieves you and a rival, and the rival's facts are more consistent with everything else the model found, you can lose the citation on a query that never mentioned the rival. You're being compared whether or not the user asked for a comparison.

What this means for how you show up

I'll be concrete about what I'd actually change, because the abstraction is only useful if it turns into work.

State claims the model can verify. "We're a leading provider" is unverifiable and gets treated as noise. "We've operated in Salt Lake City since 2009 and focus exclusively on commercial HVAC" is specific, checkable, and consistent across a reasoning trace. Specific, checkable facts survive the thinking step. Marketing adjectives don't.

Show your reasoning, not just your conclusions. If you say a particular approach is better, say why, with the mechanism. The model is building a line of reasoning; give it the raw material. Content that explains *how* something works gives the model something to reason with, and it tends to get pulled into the answer because it's load-bearing.

Get corroborated. This is the uncomfortable one, because you don't fully control it. The single strongest signal to a model that's checking consistency is that other independent sources say the same thing about you. Your own site asserting a fact is one data point. A directory, an industry association, a news mention, and a review platform all reflecting the same fact is a pattern the model can trust. Consistency across sources is the thing the reasoning step is built to detect.

Don't contradict yourself across your own surfaces. If your homepage, your About page, your Google Business Profile, and your LinkedIn give three different descriptions of what you do, a reasoning model notices, and inconsistency reads as unreliability. Pick one clear account of who you are and keep it identical everywhere. This is cheap to fix and more important than it looks.

Make one claim you can genuinely own. Reasoning models are good at detecting when everyone in a category says the same generic thing. Sameness is a reason to cite no one in particular — or to default to the biggest brand. If you can stake out one specific, defensible position or specialty, you give the model a reason to name you specifically rather than fold you into "providers like X."

The honest limits of optimizing for this

I don't want to oversell control here. You cannot see the reasoning trace, so you can't audit exactly why you were or weren't cited on a given answer. The models change their deliberation behavior with every release, sometimes significantly, and what survived the thinking step last month may not this month. Anyone selling you a guaranteed method for "ranking in reasoning models" is selling you certainty that doesn't exist.

What you can do is stop optimizing for a retrieval-only world that's already fading. The through-line across every reasoning model I've watched is that they reward sources that are specific, consistent, corroborated, and explained — and they punish vagueness and contradiction more harshly than plain search ever did. That's not a trick. It's closer to what being genuinely authoritative has always meant. The models are just getting better at telling the difference, and the extra thinking step is how they do it.

If you treat every page as something a skeptical, careful reader will interrogate before trusting — because that's now literally what happens — you'll be building for where AI search is going rather than where it was.

A quick way to pressure-test your own content

Here's a self-audit I use, and you can run it without any special tooling. Take one page you'd want cited and read it the way a reasoning model would — not "is this nicely written" but "would this survive scrutiny."

Ask it four questions. Can every important claim on this page be checked against something outside the page? If the answer is mostly no, you've written assertions, not evidence, and a reasoning model has nothing to lean on. Does the page explain mechanisms, or only announce conclusions? Conclusions without the "why" are the first thing deliberation discards. Would an independent source — a directory, a review platform, a news mention, an association — corroborate the core facts here? If nothing corroborates you, you're a single unverified data point. And finally: does anything on this page contradict what you say elsewhere? Contradiction is the cheapest way to get filtered out, and the easiest to fix.

If a page passes all four, it's the kind of source that tends to survive the thinking step. If it fails two or more, no amount of keyword work will save it, because the problem isn't that the model can't find the page — it's that the model doesn't trust it once it has. That reframing, from findability to trustworthiness under scrutiny, is the single most useful shift I can hand you. Everything else in AEO for reasoning models is a consequence of it.

Key takeaways

- Reasoning models add a hidden deliberation step between retrieval and answer, and that step is skeptical — it looks for reasons to distrust a source before citing it.
- The citation no longer goes to whoever ranks first; it goes to whoever survives the model's own scrutiny for consistency, corroboration, and defensibility.
- Specific, checkable facts survive the thinking step; marketing adjectives and vague claims get flagged as noise and dropped.
- Explaining the reasoning behind a claim gives the model raw material to fold into its own trace, so explained claims travel further than bare assertions.

- You're compared to competitors even on non-comparative queries, because the reasoning step can discard your source mid-process if a rival's facts are more consistent.
- You can't see or fully control the reasoning trace, but you can build for it: be specific, consistent across your surfaces, corroborated by others, and willing to own one clear position.

Frequently asked questions

Is this the same as ChatGPT's or Perplexity's "deep research" mode?

No. Deep research runs many searches over several minutes to assemble a report. Reasoning is a shorter, internal deliberation step that can happen on a single ordinary query — the model thinks through an answer before writing it, even when it isn't doing extended web research. Both reward trustworthy, consistent sources, but reasoning affects everyday answers, not just long research runs.

How would I even know if a reasoning model dropped my source?

You usually can't see it directly, because the reasoning trace is hidden. The practical signal is indirect: if you rank well in traditional search but still aren't named in AI answers for the same questions, that gap is often a trust-and-consistency problem rather than a visibility problem — exactly the kind of thing the deliberation step filters out.

If I can't see the reasoning, what's the single most useful thing to fix first?

Consistency across your own surfaces. Make your homepage, About page, Google Business Profile, and social profiles tell one identical story about who you are and what you do. A reasoning model that finds you contradicting yourself treats that as unreliability, and it's one of the few high-impact things you fully control.

About the author. Scott Tischler is the Founder & Chairman of AIrecommend.ai and a practitioner-authority on AI search and Answer Engine Optimization. With 20+ years in marketing technology — including American Express, MetLife, and UBS — and executive and professional study at Wharton, Harvard, and Oxford, he helps businesses become the ones AI recommends.

© 2026 Scott Tischler · scotttischler.com
scott@airecommend.ai · AIrecommend.ai