

STRATEGY

Winning the 'X vs Y' Comparison in AI Answers

By Scott Tischler · September 9, 2026 · 8 min read

Some of the highest-intent moments in any market are comparison queries — "X vs Y," "which is better," "should I use A or B." These are people at the edge of a decision, and AI answers them constantly. If you're not shaping how those comparisons get made, someone else is deciding how you stack up.

Comparison queries matter more than their volume suggests because of where they sit in the journey. Nobody asks "X vs Y" idly — they ask it when they're close to choosing and want help deciding. Win that moment and you're influencing the decision itself, not just building awareness. And AI is uniquely good at comparison answers: it can synthesize both sides into a clean, balanced-looking response in seconds. The question is whose framing and facts it's synthesizing from.

Why the framer usually wins

Here's the dynamic that most businesses miss. When an AI builds a comparison, it pulls from sources that have already laid out the comparison clearly — comparison pages, structured breakdowns, honest "here's when to choose each" content. Whoever produced the clearest, most credible comparison framework becomes the reference the model leans on. And the source that framed the comparison tends to frame it in terms where it looks strong.

This isn't about lying — a comparison that dishonestly trashes the alternative reads as biased and gets discounted. It's about being the one who defined the dimensions of comparison. If you're the source that clearly explained "here's what actually matters when choosing between these, and here's when each wins," the model adopts your framing, and your framing naturally centers the dimensions where you're genuinely better.

The counterintuitive move: compare yourself honestly

The instinct is to avoid naming competitors — why give them oxygen? For AI visibility, that instinct is backwards. The businesses that win comparison queries are usually the ones willing to write honest, specific comparisons that include their competitors by name. A fair comparison that says "choose them

if you need X, choose us if you need Y" is exactly the kind of balanced, useful content AI reaches for — and it puts you in the comparison rather than leaving you out of it.

The honesty is load-bearing here, and I want to be clear about why. A comparison that pretends you win on every dimension is transparently self-serving, and both readers and models discount it. A comparison that concedes where the alternative is genuinely better earns credibility — and that credibility is what makes the model trust your framing on the dimensions where you do win. You have to give up the fantasy of winning everything to actually win the comparison.

How to build comparison content that gets cited

- **Name the real alternatives.** Write comparisons that name the specific competitors people actually weigh you against — the exact "X vs Y" people type. Vague "us vs the other guys" content doesn't match the query.
- **Define the dimensions that matter.** Lay out the criteria a real buyer should use to choose, honestly. This framework is what the model adopts.
- **Be genuinely balanced.** Concede where the alternative wins. The credibility you earn is what makes your framing trusted on the dimensions where you win.
- **Make it extractable.** Clear structure — a criterion, then who wins it and why — so the model can lift the comparison cleanly. Buried, meandering comparisons don't extract.
- **Give a clear "choose each when."** The single most useful thing in a comparison is honest guidance on who each option is right for. That's the sentence AI loves to quote.

What to avoid

Two failure modes. The first is the hit piece — content that exists only to trash a competitor. It reads as biased, gets discounted, and can damage your credibility. The second is the fake-balanced comparison that's obviously rigged so you win everything. Both fail for the same reason: they're not honest, and AI is increasingly good at recognizing self-serving comparison content. The whole strategy depends on being the genuinely useful, fair source — because that's the source a model trusts to frame a decision.

Done right, comparison content is some of the highest-leverage AI-visibility work available. It targets people at the moment of decision, it puts you into comparisons you'd otherwise be left out of, and it lets you shape the criteria on which you're judged — all by being more honest and more useful than anyone else bothered to be. Write the comparison your customers are already asking AI to make, and make it the best one available.

Key takeaways

- ♦ Comparison queries ('X vs Y,' 'which is better') are high-intent decision moments AI answers constantly — shaping them shapes the decision.
- ♦ AI builds comparisons from sources that laid the comparison out clearly, so whoever framed it best becomes the reference the model leans on.
- ♦ The counterintuitive move is to compare yourself honestly against named competitors — it puts you in the comparison instead of leaving you out.
- ♦ Honesty is load-bearing: conceding where the alternative wins earns the credibility that makes your framing trusted where you do win.
- ♦ Build extractable comparisons: name real alternatives, define the criteria, stay balanced, and give a clear 'choose each when' guide.
- ♦ Avoid hit pieces and rigged 'balanced' comparisons — both are self-serving, get discounted, and can damage credibility.

Frequently asked questions

Should I really name my competitors in my content?

For AI visibility, usually yes. The businesses that win comparison queries write honest comparisons that name the specific competitors people weigh them against. That puts you inside the 'X vs Y' answer rather than leaving you out, and a fair, balanced comparison is exactly the kind of content AI reaches for. Avoiding competitors' names just cedes the comparison to whoever was willing to write it.

Won't an honest comparison that admits a competitor's strengths cost me sales?

It's counterintuitive, but conceding where a competitor genuinely wins is what earns credibility — and that credibility is what makes the model (and readers) trust your framing on the dimensions where you win. A comparison that pretends you win everything is transparently self-serving and gets discounted. You give up the fantasy of winning everything to actually win the decision.

What's the single most valuable element of a comparison page?

A clear, honest 'choose each option when...' section. Concrete guidance on who each choice is right for is the most useful part of any comparison and the sentence AI most loves to quote. Pair it with a criterion-by-criterion breakdown so the model can extract the comparison cleanly.

About the author. Scott Tischler is the Founder & Chairman of Alrecommend.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.