CiteHQ.io

BEST IS A FORMULA. REAL ISN'T.

On AI, citation, and why exposure stops standing in for merit.

By David Maruyama

THE VISION — ESSAY 4 OF 6

This is the fourth in a six-part series — a personal reflection on what AI actually changes, published on Tuesday nights over the coming weeks. There's no pitch here and nothing for sale: just honest reflections on AI, and an attempt to actually act on them, in public, alongside anyone else figuring out the same thing in real time. Coming next: What a Beacon Actually Is.

Ask most people why they chose the last vendor, tool, or brand they picked over an alternative, and you'll usually get some version of the same answer: it was the best one. Ask a follow-up question — best by what measure, compared against what — and the confidence tends to drop fast. Most of the time, "best" wasn't measured. It was just the option that had reached them the most, from the most directions, the most times.

That's worth being honest about, because it's not really a judgment of quality. It's a measurement of exposure. Whoever spent the most to be in front of us, whoever got there first and locked in the advantage of being the familiar name, whoever showed up in the most places — that's usually who wins the word "best," regardless of whether they deserved it. In marketing, "best" often just means "the obvious choice," not necessarily the right one. Our minds quietly treat visibility as a stand-in for merit, because that's a much easier judgment to make than evaluating the thing itself. Best, in other words, is a formula: spend, incumbency and exposure, added up.

I think there's a better word for what we're actually looking for, and it isn't "best." It's real.

"Best" always needs a comparison to mean anything — best compared to what, ranked by whom. "Real" doesn't need a competitor in the room at all. Something either holds up against reality, or it doesn't. It doesn't need to beat an alternative to be real; it just needs to be what it claims to be, under whatever scrutiny gets pointed at it.

For a long time, most businesses didn't have to worry much about that distinction, because the scrutiny wasn't there yet. An advantage built on being

first, on inherited credibility, on a gap in information nobody else could access, could last for years without ever being tested. Arbitrage. Acquired position. A head start that had nothing to do with whether today's product or service was good. As long as nobody could see clearly enough to check, the story could hold. That is changing, and not by a small margin. When information becomes this easy to capture and this precise to process, an advantage that depends on staying unverified doesn't survive. Familiarity stops being enough of a moat, because the tools to check the substance behind it are now as available as the familiarity itself.

For most of business history, a brand had two sides, managed almost independently: the storefront, and everything holding it up. How a brand looked to human eyes could differ from what happened behind it, and nobody was checking one against the other, so nothing forced them to match. AI changes that. It can read across far more than any person could — what the brand says about itself, what customers say, what outside sources say — and see whether the image and the substance describe the same thing. The formula that produced "best" can now be taken apart. I don't think consistency on both sides is optional anymore.

Take a company that spends heavily on a brand voice built around care — warm language, a promise of always putting the customer first, the kind of image that used to work simply because it was repeated often enough and looked good doing it. Before, if you asked a person how that company treated its customers, the answer mostly came from what the company had said about itself, reinforced by how polished it looked saying it. Now ask an AI system the same question, and it isn't reading only the brand voice — it's synthesizing across public reviews, complaints, forum threads and press coverage, and when the real pattern doesn't match the promise, it can say so. The gap that used to be invisible, protected by the fact that nobody could check both sides at once, is now just the answer to a simple question.

There's already evidence of this playing out, not just a hopeful theory. In a 2025 Ahrefs study of 15,000 queries, only about 12% of the links cited by ChatGPT, Gemini and Copilot ranked in Google's top ten for the same prompt, and about 80% didn't rank anywhere in Google. That's a strange number if AI citation were just old-fashioned ranking wearing a new coat. It varies by engine, and it doesn't prove AI rewards substance. These systems may simply favor different kinds of sources. But it does suggest the incumbency shortcut doesn't carry over automatically, at least for now.

I want to be careful not to oversell that, though. It isn't a permanent guarantee, and it isn't automatic. There's already a real effort underway to manufacture the appearance of "real" for AI the same way an entire industry once manufactured the appearance of relevance for search engines — synthetic authority, content built to be cited rather than to be true. Being real doesn't win on its own. It has to be defended, the same way anything real has always had to be defended against something convincingly fake standing next to it.

That leaves the harder question, and I don't have a clean answer yet: what happens to a brand once an inconsistency is caught, not just suspected — and what does it take, with AI doing the checking, to close the gap instead of managing around it?

So here's what I'd ask you: has there been a moment — in your own work, your own industry, a purchase you made or a decision you watched someone else make — where the real thing won out over something that was simply better marketed? I'd like to hear the specific story, not just the principle. That kind of evidence would convince me this shift is real more than any research I could quote.

Source: Ahrefs, “Only 12% of AI Cited URLs Rank in Google’s Top 10 for the Original Prompt” (Aug 11, 2025). ahrefs.com/blog/ai-search-overlap

---



Connect with us: feedback@citehq.io