NOTES FROM THE BOT LAB · EDITION OF NOVEMBER 25, 2026

You Made the Shortlist on the Desk — and Still Got Left Out of the Sentence

By Will Tygart · Tygart Media · Canonical edition at https://tygartmedia.com/shortlist-on-the-desk-left-out-of-the-sentence/

It is one thing when nobody invites your company to bid on a job. You shrug, figure they called somebody across town, and go back to work. It stings a whole lot differently when your bid folder is sitting face-up right on top of the buyer's desk in the conference room — and when the meeting starts, they read somebody else's summary out loud.

Most conversations about AI search treat visibility like a coin with two sides: either the bot found you or you are invisible. Today, looking inside the raw job payloads of ten different AI engines on our own site, we watched the middle state happen over and over again. We call it retrieved, not cited — making the shortlist on the desk and still getting left out of the sentence.

Telling This One on Ourselves First

I would rather tell an uncomfortable number on my own shop than preach from a balcony. Earlier today we ran ten core topics from tygartmedia.com across ten AI engines — ninety-eight settled runs in total. On Bing Copilot's web search layer, our article ranked #1 five times and #2 once. Six times out of ten, Copilot reached into its search index, grabbed our page right at the top of the stack, set it on the workbench — and withheld the inline footnote on five of them.

Across our wider 295-run B2B benchmark yesterday, that exact pattern held steady: Bing Copilot retrieved the target domain in 74.0 percent of runs and cited it inline in 6.0 percent. Even Claude Sonnet did its own version today on a software pricing test: its raw tool log pulled seventeen candidate web pages into background memory, then cited only seven unique URLs in the final prose. Ten pages made the room and missed the microphone.

What 91 Head-to-Head Losses Actually Taught Us

When you lose a slot you already ranked #1 for, guessing why is dangerous. On our first pass comparing our skipped pages against the winners, it looked like rival articles had twice as many bulleted lists. Then we realized our first script was accidentally counting rival mega-menus and footer links. Once we stripped navigation, headers, and footers cleanly off 135 rival pages and PDFs — isolating 91 like-for-like editorial HTML duels — the bullet-list gap vanished completely: 19.0 versus 18.0 items, pure statistical noise.

In fact, our skipped pages already beat the cited rivals on structured tables and quantified numbers in the first three hundred words. Across every content metric we tested, peer rivals beat us in exactly one place: the length of the very first paragraph. Cited winners opened with a self-contained 47-word synthesis paragraph versus our clipped 31-word intros, winning 61.5 percent of paired matchups. Pair that with 3,413 inline citation anchors showing engines drop links at more than double the rate in the opening quarter of a page, and the culprit was staring right at us. Even our own Markdown twins had been squeezing intros down to seventeen-word stubs.

Stop Mumbling at the Doorway

We wrote in Cited Is Not Chosen that getting quoted is not the same thing as getting hired, and in We Come From Water Damage that most invisibility starts as structural plumbing. Today added the quieter lesson in between. On one crawl-speed topic today where our page carried a full opening answer, a clean table, a companion PDF, and a synced Markdown twin — the exact formats we tested in our /llms.txt and Markdown twin log study — Perplexity cited our .md file at #1 and Microsoft's enterprise Copilot cited our PDF directly.

Once you respect the 7-to-14-day curing window and account for why big incumbents get away with locked doors, check how your best pages open. If an engine can lift your first forty-five to fifty-five words straight onto the conference table and answer the buyer's whole question cleanly, you stop getting left in the folder.

Measured Signal (91 Peer HTML Duels)	Cited Rival Median	Skipped Page Median	Statistical Verdict & Meaning
Opening synthesis paragraph (first_para_words)	47.0 words	31.0 words	Rivals win 61.5% of pairs (p = 0.0116) — the #1 real prose gap
Inline citation anchor position (3,413 spans)	4.12% in top 25%	1.97% downstream	2.1x concentration near the top of the page (p = 0.00137)
Comparison tables on page (table_count)	0.0 tables	1.0 table	Our pages already led (p = 1e-6) — tables alone need a lead synthesis
Quantified numbers in first 300 words	7.0 numbers	13.0 numbers	Our pages already led (p = 0.00078) — raw stats need full sentences
Body bullet items (after stripping nav/footer)	19.0 items	18.0 items	No difference (p = 0.275) — naive scrapers were counting menu links

Key Questions Answered: What does "retrieved, not cited" mean in AI search? It means an AI assistant pulled your URL into its background search results or tool scratchpad as a top candidate, then synthesized its final answer quoting rival sources without linking to you. **Why did rival pages beat higher-ranking pages with more tables in your 91-pair autopsy?** After stripping navigation chrome, cited peer pages won consistently on opening paragraph completeness — averaging 47 words of self-contained synthesis in paragraph one versus 31 words on skipped pages. **Do AI engines ever cite PDFs or Markdown (.md) files directly instead of HTML?** Yes. In our 10-engine baseline on tygartmedia.com, Perplexity cited a companion .md twin at rank #1 and Copilot cited a 2-page reference PDF directly alongside the HTML article.