

How ChatGPT became our #1 referrer

The CitedStack playbook

A solo builder's account of an affiliate site that gets more buyers from ChatGPT than from Google Search — off zero ad spend, at near-zero domain authority — and the discipline that keeps it there.

1. The discovery

I did not set out to optimize for ChatGPT. I opened Google Analytics one week to figure out why a small buying-guide site was earning at all, given it ranks on page 2 for almost everything and has effectively no backlinks.

The channel report was not what I expected. The top referrer was not Google organic. It was **ChatGPT**. Not close, either — ChatGPT sent more buyers than classic Google search did. Add Google's AI Overviews and Perplexity, and the picture was blunt: **AI answer engines were the business. Classic blue-link search was a rounding error (~5%). Pinterest, which I had spent weeks on, sent nothing that converted. Google Discover: zero.**

That reframed everything. The site was not winning search. It was winning *citations* — getting quoted inside AI answers, with a link, at the exact moment someone was deciding what to buy.

2. Why this happens (and why it is good news if you have no authority)

Traditional SEO is a popularity contest you lose for two years while you build links. AI citation is not. Answer engines assemble a response by pulling *passages* from pages they judge relevant, well-structured, and specific — and they draw from a far wider pool than the top 3 blue links. A page on a no-authority domain can get cited if it answers the question more cleanly than the page that outranks it.

Three things make this true:

1. **Engines cite passages, not pages.** If your answer is a clean, self-contained 40-to-60-word block, it can be lifted verbatim. If it is buried in setup prose, it cannot, and the engine quotes whoever front-loaded it.

2. **Structure and specificity beat authority at the margin.** A comparison table with real criteria, a sourced statistic, a dated "best X under ₹Y" verdict — these are more liftable than a higher-DR page of marketing fluff.
3. **The Princeton GEO research quantified it** (Perplexity, KDD 2024): citing sources (+40%), adding statistics (+37%), and quotations (+30%) are the biggest levers — and *low-authority sites benefit more*, up to a ~115% lift from citations. Keyword stuffing, meanwhile, actively *reduces* AI visibility (-10%).

The uncomfortable, liberating conclusion: **you do not need authority to get cited. You need extractable, specific, verifiable answers.** That is a thing a solo builder with Claude Code can produce at scale. Backlinks are not.

3. The pages that actually get cited

Not all content is equal in an AI answer. On our sites, citations concentrate on a few formats, in this rough order:

- **"Best X for Y" and "Best X under [price]"** — high buying intent, and the engine wants a structured, ranked answer it can quote. This is the workhorse.
- **"X vs Y" comparisons** — comparison content wins roughly a third of AI citations because a table is the ideal liftable object.
- **"How to [solve the problem X fixes]"** — procedural, step-structured, HowTo schema.
- **Original data / cost calculators** — a page with your *own* numbers is the most citable thing you can make, because there is no more authoritative source for your data than you.

Thin, generic, undated blog posts get cited by nobody. Gated content is invisible.

4. The structure that gets lifted

Every cited page on our sites follows the same skeleton:

1. **H1 = the query, phrased how a human asks it.**
2. **An at-a-glance answer block** directly under it: the direct answer in 40 to 60 words, written to survive being copied out with zero context. This one block does most of the work.
3. **A clear verdict** ("our pick: ..., because ...").
4. **A comparison table** whose columns are the decision criteria.

5. **Sourced statistics** — every number has a named source and a date.
6. **A real FAQ** — questions taken from People-Also-Ask and search logs, each answered in a self-contained paragraph, backed by FAQ schema.
7. **Matching JSON-LD** (Article + FAQPage + ItemList/HowTo) that mirrors the visible page exactly.
8. **Freshness and attribution** — a visible "Last updated" date and a real author.
9. **llms.txt** and **robots.txt** that let the crawlers in and hand them a map.

The `aeo-blocks` skill in this pack generates all of it. But structure is only half the system.

5. The half nobody talks about: getting cited is worthless if the claim is a lie

Here is the failure mode that will quietly kill an AI-content site, and it is the reason this pack exists.

When you generate content at scale with an LLM, it fabricates. Confidently. Constantly. Not "occasionally wrong" — **systematically** wrong on exactly the facts that matter on a money page. In a single verification sweep across our sites, the numbers were brutal:

- **82% of AI-generated product links were fake** before verification — dead, wrong item, or a search page dressed up as a product link.
- In one luxury category, **23 of 29 product links pointed to the wrong item.**
- The first-pass writer approved every one of them. It had no idea.

Why this is fatal and not just embarrassing: one fabricated product, wrong price, or phantom link on a money page and you lose trust with all three parties at once — the **search/answer engine** stops citing an unreliable source, the **affiliate program** (Amazon et al.) treats a wrong displayed price as a policy violation, and the **reader** who clicks a dead link never comes back. Volume without verification is not an asset. It is a liability that compounds.

AI content sites do not die from bad writing. They die from lying.

So the second half of the system is an adversarial verifier — the `factcheck-verifier` agent in this pack — that runs before every publish and treats the draft as guilty until proven innocent. It re-checks every product, price, spec, certification, statistic, and link against a real external source, and returns a per-claim BLOCK/PASS. Nothing ships with an open BLOCK. It is the pass that assumes the first pass lied, and it is the single most important thing protecting the revenue.

6. The exact workflow

For every page, in order:

1. **Pick one query** a real buyer types, phrased naturally.
2. **Access check** — confirm GPTBot, PerplexityBot, ClaudeBot, Google-Extended, Bingbot are allowed and the page is server-rendered. (Skip this and nothing else matters.)
3. **Generate the blocks** with `aeo-blocks` — at-a-glance answer, verdict, table, sourced stats, FAQ, HowTo if procedural.
4. **Add matching JSON-LD** — Article + the block-specific schema.
5. **Verify with `factcheck-verifier`** — every claim BLOCK/PASS. Fix all BLOCKs. This gate is non-negotiable.
6. **Emit machine-readable files** — refresh `llms.txt`; add `pricing.md` for product pages.
7. **Ship, then submit** — sitemap + IndexNow/Bing so the crawlers find it fast.
8. **Monitor** — watch the GA4 channel/referrer report and, monthly, run your top 20 queries through ChatGPT, Perplexity, and Google to see where you are cited and where a competitor is cited instead. The gaps are your next pages.

7. What did *not* work (so you can skip it)

Being honest, because the receipts cut both ways:

- **Pinterest, X, Bluesky, social in general** — near-zero converting traffic for this kind of content. Time sink. Social is not the lever.
- **Chasing backlinks early** — high effort, slow, and unnecessary to *start* getting cited. Authority helps eventually; it is not the entry ticket.
- **Publishing volume without verification** — actively harmful. Fabrications accumulate and erode the exact trust that citation depends on.
- **Google Discover / classic ranking as the plan** — for a new site, this is a two-year bet. AI citation pays out in weeks.

8. The honest limits

This is a small business: an affiliate run-rate in the low tens of thousands of rupees a month, growing, off zero ad spend. The point of the receipts is not the size of the number. It is the *mechanism*: a no-authority site, using Claude Code, producing verified, extractable pages,

and getting cited by the engines where buying decisions now happen. That mechanism is what this pack hands you. What you point it at, and how far you push it, is up to you.

Questions: 31nathbikram@gmail.com. This playbook ships with the `aeo-blocks` skill and the `factcheck-verifier` agent that implement it.