Poisoned search results: how to catch them

A poisoned page has the right structure and the wrong content: an AI Overview, ten links and ads, all present, with the top results swapped. Every structural check passes. Catch it by comparing with something you trust.

Signs a page may be poisoned

- Most of the top domains changed since the last time you fetched the same query.
- Big, well-known sites for the topic are missing; small or unexpected ones lead.
- The AI Overview cites the same unexpected sources as the links.
- Channels disagree about the same query at the same time (a sign, not a verdict: check each one).
- Brand queries look fine while commercial queries look odd.

How to catch it, best to weakest

1 **Compare with its own history** — `best, needs history`

Rankings drift a domain or two a day. A sudden swap of most of the top domains is the alarm: don't trust that page.

2 **A reference page, through the same channel** — `reliable`

Keep a reference page: a known commercial query and its correct results, refreshed regularly. On every request, send the reference query in parallel through the same channel (IP, cookies, browser bundle) and compare. Match: the channel is clean, keep the answer. No match: the answer is poisoned too.

3 **Relevance check** — `weakest`

Are the query's words in the titles? Catches off-topic poison only; on-topic poison passes.

- [ ] It's not only Google: any site can serve a bot a normal-looking page with the wrong data. Apply the same checks to every source you scrape.