


Entity records fixed, and the bare name still fails.

A pre-registered intervention on a site we own. We corrected and added the records that describe the company, then asked for the bare name on the dates we had set in advance.



O

of **654** searches on topic and buying questions named us after the fix. Asked for the bare name, zero of four, three times. The one API answer that cited us cited our first paper.

REGISTERED FIRST, READ TWICE

The prediction, its falsifier and both reading dates were written down before any record changed. Both readings came back null.

Seven weeks after we fixed the records that describe Cortexa Consulting, typing our bare name into ChatGPT still finds 23 other companies and not us. We had predicted the opposite, in writing, before we changed anything. The prediction was that answer engines decide whether a company is a candidate answer from the structured records that describe it: directory profiles, business registries, a Wikidata item. If that were true, cleaning up those records should make an engine find us when someone types our bare name.

We wrote down the test, the date it would be read, and what would count as failure. Then we fixed the records and read the result twice, as scheduled. Both reads were null. In the same seven weeks, one answer to the bare name did cite a page of ours. It was our first published paper on this subject, and it was cited for what it says about the crowded name.

CORTEXA'S READ

A null is worth publishing when the test could have come out the other way. This one could have. The records were in place, the engine searched every time, and the outcome we would have reported as a win was written down in advance. It didn't happen. The only thing of ours an engine picked for the bare name was research we published, which says something about where the effort pays.



Our listing, corrected and polished. The slip in the visitor's hand still matches every line on the board.

What did we predict, and when?

On 25 August, before any record changed, we filed a hypothesis in our findings ledger: entity records, rather than content, decide whether a company is in the set an engine chooses from. The main test was four cold asks of the bare name in ChatGPT, each in a fresh session. Success meant any page of ours cited, or our name appearing when we followed up with “Are there others?”. The baseline that night was zero of four.

Two more checks were registered with it. Our daily API runner already asks the bare name 24 times a day, and any citation there would count as a signal. The other seven questions that name the company were the control: they were already cited about nine times in ten, and if they rose too, something other than the records had changed.

The first read was booked for 15 September. Two days after registering, we learned that the record with the widest reach would not take effect until the week before that read. So on 27 August, with no result in hand, we booked a second read for 15 October and wrote down why: a record that has had five days to spread has not been tested. The 30 days between the two reads, 16 September to 15 October, are what this paper calls the second observation window. We also wrote down what would spoil the second read: any further change to our records between the two dates had to be logged with its date. Two small ones were made and not logged, and we report them below.

On 18 September we registered a second prediction, this one about mechanism. Our first paper had found that the zero on topic questions is a retrieval failure. When an engine searches about a subject, its search phrases never name us, so our pages never enter the pool it picks from. Entity records describe who we are once found. They cannot change what an engine types into its own search box. So we predicted that the non-brand search phrases would still name us zero percent of the time, and that the rate at which engines search at all would stay within five points of where it was.

What did we change?

What we did is the standard advice. Practitioner guides to Generative Engine Optimization (GEO) recommend creating or correcting a Wikidata item, linking your profiles to it, and keeping the description of the company identical everywhere it appears.¹ Some go further and treat consistent business profiles as one of the first things to fix.² Most of what we did was correcting records that already existed and were saying the wrong thing. Some of it was new.

- **Location, corrected everywhere.** Our LinkedIn page listed three different cities. Our Google Business Profile was worse, and the reason is worth knowing. We had no public address, and for customer locations we had picked several large metropolitan areas across the country, the places we expected to work.

Google places a business like that at the center of the areas it serves, and the center of a nationwide list of cities is the middle of the country, so the profile put us in south-central Kansas. All five profiles we control now agree on West End, North Carolina, and the Google profile carries our exact address, hidden from public view.

- **A Clutch profile, created 25 August** and filled to the limit of the free tier. Engines began citing it the next day, which our first paper reported.
- **A Wikidata item, created 26 August**, with the legal name, founding date, headquarters, official website and a state registry identifier, most of them backed by references.
- **The Dun & Bradstreet record** gained a website field it had never carried. That record feeds more than eighty downstream business sites, so it was the change we expected to reach furthest. It took effect on 10 September.
- **Our own site's record of the company**, the structured data search engines read, gained the D-U-N-S and state file numbers on 25 August and a link to each new profile.

The record most likely to matter was the slowest to arrive. That is why the second read exists, and why it was scheduled five weeks after that record went live rather than five days.

02

CORTEXA FIELD STUDIES N° 02

What the engines did

What happened when we asked for the bare name?

Nothing, three times. Each read asked the same four questions in fresh, signed-out sessions: “Cortexa”, “Cortexa company” twice, and “What is Cortexa?”, each followed by “Are there others?”.

- **25 August, before any change:** zero of four.
- **15 September, the Dun & Bradstreet field five days old:** zero of four.
- **16 October, the same field 36 days old:** zero of four, and absent from all six follow-up lists. Across ten turns of conversation, no page of ours appeared once.

Every test we registered, and what it returned

Asked cold in the app, the bare name stayed shut, and it opened once through the API. The full name worked every time. Searches on questions that do not name us never named us.

Bare name, asked cold in the app, three reads

0 of 12

Bare name through the API, second window

1 of 552 · 0.2%

The one answer that cited us, 7 October: our first paper, Cortexa Field Studies N° 01

Full name, signed out in the app

5 of 5 · 100%

Searches on non-brand questions that name us

0 of 654

Searches on brand questions that name us

1,065 of 1,075 · 99.1%

The other seven brand questions, API, second window

3,547 of 3,864 · 91.8%

Cortexa Field Studies N° 02 · second window 2026-09-16 to 2026-10-15 · cold reads on 25 August, 15 September and 16 October · the one API hit (7 October) cites the N° 01 paper · searches from the two engines that expose them

Every test we registered, side by side. The bare name opened once, through the API, and the page it opened on was our own first paper.

One thing did change, and it makes the null harder rather than softer. In August the engine needed two nudges before it would search for a bare “Cortexa”. In September it needed one. In October it searched on the first turn of all four asks. So the engine is looking more readily each time, and it is still not finding us when it looks.

What it finds instead is a crowd. The October lists named 23 distinct companies called Cortexa, about the same number as in September. The size held steady while the membership moved: six of the 23 had never appeared in our records before. One of them is an AI strategy advisory for leadership teams, which is close to our own description of ourselves under another owner. None of these companies is doing anything wrong. They hold the name too, and some held it first.

Did anything of ours get picked for the bare name?

Once. Our daily runner asks the bare name 24 times a day through OpenAI's API, in three wordings. From 7 August to the first read it had produced no citation in 680 tries. In the second window it produced one in 552. A second runner asks it once a day on each of three engines, and that one stayed at zero over the same weeks, 81 asks in all.

It came on 7 October, to the question “What is Cortexa?”. The answer listed five other companies, among them a therapy-practice analytics firm, an AI medical scribe and a British learning-technology company, and asked which one we meant. Then it closed by noting that the name is used by many unrelated organizations, and cited a page for that line. The page was Cortexa Field Studies N° 01, the first thing we published on this subject.

So the engine did not decide that we are one of the Cortexas. It found our measurement of the crowded namespace and used it as evidence that the namespace is crowded. That is a different kind of citation from the one the records were meant to earn. The records were supposed to make the bare name resolve to the company. What happened instead is that a piece of research was retrieved and quoted for what it says.

Our server logs show how it got there. OpenAI's search crawler fetched that paper on six days during the second window, twice on 6 October, the last time at 20:00 UTC. The answer that cited it came from the run that started at 07:05 UTC the next morning, about eleven hours later. ChatGPT's live fetcher read six other pages of ours in that hour, but not the paper, so the engine cited it from what it had already read. The same crawler had read the page on five earlier days without a citation following any of them. A recent crawl came before the one citation, and crawling alone was not enough to produce one.

CORTEXA'S READ


One citation in 552 is an observation, and we will not call it a trend. But it is the first time in this study that an engine chose a page of ours for a question that does not give our full name, and the page it chose was published research. Seven weeks of record fixes did not do that.

Did the records change anything else?

No. The registered control held: the seven other questions that name the company were cited 91.3 percent of the time between 20 August and the first read, and 91.8 percent during the second window. Day by day, neither line moved when the records changed on 25 August or when the widest one went live on 10 September.

Nothing stepped when the records landed

Daily citation rate through the API, from 20 August to the end of the second window. The records changed on 25 August and the widest one took effect on 10 September. Neither line moved.



The records landed on 25 August and the widest one went live on 10 September. Neither line stepped. The one point off zero on the bare-name line is the N° 01 citation.

Did the records change what the engine searches for?

No, which is what we predicted. Between 16 September and 15 October our daily runner asked the same 24 frozen questions it has asked since 7 August, and two of the engines report the search phrases they generate. We counted every phrase they produced for the questions that do not name us.

- **654 search phrases, and none of them named us.** Zero on both engines, the same as the frozen baseline from August.
- **The rate of searching barely moved.** OpenAI searched on 17.9 percent of non-brand questions against 15.6 before. Gemini searched on 3.0 percent against 7.0.
- **Brand questions still worked as before.** On questions that name the company, 98 percent of OpenAI's search phrases and all of Gemini's named us.
- **Topic citations stayed at zero** on all three engines, every day of the second window: none in 1,296 answers to topic and buying questions, and none in 845 more from two smaller question sets that also

never name us.

How often the engines searched on questions that do not name us

Registered in advance: within five points of the baseline. Both engines stayed inside it. Gemini's drop follows Google's own change of model.



Cortexa Field Studies N° 02 · wave arm, topic and buying-intent questions · in the second window Gemini ran entirely on a newer model than for most of its baseline

Both engines stayed inside the five-point band we registered. Gemini's drop came from Google's own change of model.

Gemini's four-point drop needs a note, because it is the larger move and it has nothing to do with us. Google changed the model behind that endpoint in mid-August, and the newer model searches on almost nothing outside one subject. Every day of the second window ran on the newer model, while the August baseline straddles the change. We registered that risk before the second window opened.

We also left room for a result we did not expect and would have liked: topic citations moving off zero while the search phrases still never named us. That would have meant we were getting chosen without being searched for. No topic citation appeared in the whole second window.

Were we being read at all?

Yes, steadily. This is the half of the study most measurement cannot see, and owning the site is what lets us see it. Across the 30 days of the second window our edge logged 7,961 requests from named AI crawlers. Those names are self-declared, and our edge could not confirm about one request in ten, so we treat the count as a ceiling rather than a measure of attention. The paper that was cited on 7 October had been fetched by OpenAI's search crawler 33 times, counting its figures, before and after that day.

One observation from the October read we cannot pin down, and we report it for that reason. In the hour that held the first cold ask, ChatGPT's live-fetch agent requested five of our pages: the home page, About, the AI and Studio service pages, and Contact. None of the four answers cited us. Our crawl logs are grouped by the hour, so we cannot tell whether those fetches came from our session or from someone else asking about us in the same hour.

So does the engine know us at all?

It does, by our full name. That was never the open question, but it matters for reading the null. On 16 September, signed out and on a phone, we asked four full-name questions such as “Who is Cortexa Consulting?”. All four cited our site. On 16 October, minutes after the third cold read and on the same phone, we asked “Cortexa consulting” once more. It cited us on the first turn, with an accurate description of what we do, where we are and who founded the company.

When we tapped the app's own follow-up about legitimacy, the answer cited our site, our Clutch profile and the founder's LinkedIn page. It also named a dissolved British company with nearly our exact name and set it aside as a separate historical record. On 7 August that same British company is what the engine answered with when it failed to find us. Telling the two apart has now happened in two of five full-name sessions.

We cannot credit the records with that, and we will not. The full name was already cited before the intervention, and two sessions are not a rate. What the comparison does show is where the problem sits. With the full name the engine resolves us. With the bare name it has 23 candidates, and the search results for “Cortexa” belong to whoever ranks for “Cortexa”.

One more check closes a door. A model can answer from memory without searching, and if a brand answer named a true fact about us with no search behind it, the company would be in the model's training data. During the second window OpenAI answered 54 brand questions without searching, and Gemini three. None of those answers named a single specific fact about us.

Isn't pre-registration common in GEO now?

More than it was when we started, and we want to claim only what is ours. In August we could find almost no pre-registered intervention studies in GEO. By late September there were at least two serious efforts. The GEO Lab keeps a public, dated log of its experiments and instruments its own server logs.³ It has also published a pre-registered test of whether agent-readiness scores predict AI citation across 36 sites, deposited in advance, and reported a null.⁴

So the claim this paper makes is narrower: a pre-registered intervention on a site we own, measured on both sides, from a cold start. We can see the answers and the server logs, we started from a domain that did not exist in June, and the test was written down before the change. We know of no other study that combines all of those, and we expect that sentence to age.

What did we get wrong?

- **We kept publishing.** The registration said the test needed no content changes, so the two mechanisms could be told apart. Between 25 August and the second read we published twelve Ground Truth articles and N° 01 itself, and the one bare-name citation landed on N° 01. So content and records both changed during the window, and the content is the likelier explanation for that one hit.
- **Two record changes went unlogged.** On 26 September the founder's identifier was added to our site's record of the company. Late on the evening of 14 October, about 01:00 UTC on the 15th, we created a YouTube channel, which still holds no videos, and that record links to it. A new profile is the kind of change the window was meant to exclude. It came a week after the one citation and about a day and a half before the second read, too late to plausibly move either, and the read was null anyway. Neither change was written into the ledger at the time.
- **The conditions moved between reads.** The August baseline ran signed in, in a temporary chat, before we learned that a temporary chat could still reach the account's memory. Both later reads ran signed out. September was a laptop on home WiFi in North Carolina; October was a phone on a cellular network in Michigan. All three reads are zero, so none of these differences can have hidden a positive, but a positive would have needed them controlled.
- **We never wrote down where the second window started.** The registration said to read from “the block start” and did not define it. We chose 16 September, the day after the first read, and state it in the ledger. Reading from 22 September instead changes no conclusion.
- **Two gaps in collection.** The API account ran out of credit from 11 to 21 September, and nothing at all was collected on 30 September. Neither gap is filled in or averaged over.

- **We changed the site the same morning.** A separate study's on-page changes shipped about thirty minutes after the October asks. The read came first, so it cannot have been affected, but the timing was closer than we would choose.

What would we tell a client?

- **Fix your records for accuracy, and expect no more than that.** Correct, consistent records are hygiene. In our test they did not make a contested name findable and did not touch topic questions at all.
- **Own the full name in everything you publish.** In every signed-out session where the engine was given “Cortexa Consulting”, it found us. Given “Cortexa” on its own in a fresh session, the app never has.
- **Publish something worth quoting about your own corner.** The one time the bare name led anywhere of ours, it led to research we had published. One case is thin, and it is still the only door that opened.
- **Check the namespace before you name anything.** A short brand word shared with two dozen companies has a cost you can measure, and no record we could file paid it down in seven weeks.
- **Ask any vendor selling entity work what it will move, and by when.** If they cannot say what result would show it failed, they are not making a claim you can check.

What this study cannot tell you

Four asks per read is a small number. It can detect a change from never to often, and that was the change the hypothesis promised. It cannot detect a slow drift. The cold reads come from one consumer app, signed out, so they say nothing about how a signed-in user with history would be answered.

The registration named one way the test could be void rather than negative: if the Wikidata item were deleted mid-window. It was still live when we checked on 19 October. Seven weeks may be too short for every downstream directory to pick up a change, and we cannot see which ones did. The company also aged from seven weeks to fourteen during the test, which the registration named as a confound we cannot separate from the records. We also cannot separate the two halves of what we changed. Most of the intervention corrected records that contradicted each other; some of it added new ones. If the needle had moved, we could not have said which half moved it. It did not move, so the question stays open in both directions.

The one bare-name citation cannot be traced to a single cause either. The N° 01 paper had been live for two weeks and read by the search crawler several times, and the records had been in place for six. We think the content explains it, because the answer cited the paper for exactly what the paper says, but one case does not let us rule the records out.

And this is one company with a contested name. A brand with a word nobody else uses would very likely see a different result, because its bare name and its full name would be the same question. Ours are not.

Sources

1. Frase — “Entity Optimization for GEO: The 2026 Practitioner Guide”

<https://www.frase.io/blog/entity-optimization-for-geo>

Vendor guide, cited as an example of the advice this study tested. Described from search results on 2026-10-19; the page could not be fetched from our research environment that day.

2. SkySync — “The GEO Guide: Generative Engine Optimization & AEO in 2026”

<https://www.skysync.nyc/geo/guide>

Agency guide, cited as an example of the advice this study tested. Described from search results on 2026-10-19; not fetched.

3. The GEO Lab — The GEO Log: research updates and experiments

<https://thegeolab.net/log/>

Practitioner research log, not peer-reviewed. Described from search results on 2026-10-19; the page could not be fetched from our research environment that day.

4. Ferreira (The GEO Lab) — “Agent-Readiness: Does Conformance Predict AI Citation?” · pre-registered, n = 36, null result; deposited at Zenodo, DOI 10.5281/zenodo.20788222

<https://thegeolab.net/prereg-agent-readiness-citation/>

Self-published and pre-registered; the deposited files were not retrieved by us. Figures as reported by the author.

Appendix · how every figure was captured

Most of the numbers in this paper are our own readings rather than someone else's published figures, and on the page the two look identical. This is the difference. Below are the questions we asked and when we froze them, every headline number beside the command that reproduces it, what each collection day looked like, and the counting decisions the numbers rest on.

1 · The questions, fixed before anything was measured

Written and frozen before the first reading, and not edited since. A query set you tune while measuring is a query set that measures your tuning. The checksums are there so you can tell whether the set below is the one that produced the numbers.

Question set	24 questions, frozen 2026-08-07 · sha256 <code>f2cd9687db2d44b58930d3a690c30c87...</code>
Wording variants	3 per question, frozen 2026-08-09 · sha256 <code>f7611854b31af3e7eaa97fcfd03a6213...</code>
Measured	2026-09-16 to 2026-10-15 · the engine changed on 2026-08-20, and nothing is pooled across it

ID	ASKS ABOUT	THE QUESTION, WORD FOR WORD
D01	Names the company	Who is Cortexa Consulting?
D02	Names the company	What does Cortexa Consulting do?
D03	Names the company	Cortexa
D04	Names the company	Cortexa Consulting services
D05	Names the company	Is Cortexa Consulting a marketing agency or a technology consultancy?
D06	Names the company	Who founded Cortexa Consulting?
D07	Names the company	Cortexa Consulting North Carolina
D08	Names the company	What is Cortexa Ground Truth?
T01	Topic	What does agentic AI actually mean?
T02	Topic	Why does AI make things up?
T03	Topic	What is RAG in simple terms?
T04	Topic	Should we fine-tune a model or use RAG?
T05	Topic	What is a token in AI and why does it cost money?
T06	Topic	Is prompt engineering a real skill?
T07	Topic	Why do AI demos fail in production?
T08	Topic	What are evals in AI?
T09	Topic	What does multimodal AI mean for marketing?
T10	Topic	Which AI agents actually work in production?
T11	Topic	What is MCP, the Model Context Protocol?
T12	Topic	What is a vector database and why would I need one?
C01	Buying intent	AI consulting for marketing agencies
C02	Buying intent	market research for technology clients
C03	Buying intent	white label AI consulting for agencies
C04	Buying intent	who does AI market research for tech brands

2 · Every figure, and the command that reproduces it

Run the command against the filed readings and you should get the figure beside it. The second line under each command is the decision that number rests on, which is the part no tool can tell you.

THE FIGURE IN THE PAPER	HOW TO REPRODUCE IT
0 of 4 · 0 of 4 · 0 of 4 the bare name, asked cold in the app on 25 August, 15 September and 16 October	<pre>node tools/geo-probe/cold-reads.mjs</pre> Four fixed asks, fresh signed-out sessions (the August baseline ran signed in). Success, defined in advance: any page of ours cited, or our name in an “Are there others?” list. Files: data/geo-probe/manual/*-d03cold*.
0 of 654 search phrases on non-brand questions that name us	<pre>node tools/geo-probe/fanout-report.mjs --from=2026-09-16 --to=2026-10-15</pre> Wave arm, OpenAI and Gemini, the two engines that expose their searches. Topic and buying-intent questions. Baseline 0 of 853 from N° 01's window.
17.9% · 3.0% how often OpenAI and Gemini searched on non-brand questions in the second window	<pre>node tools/geo-probe/fanout-report.mjs --from=2026-09-16 --to=2026-10-15</pre> Registered band: within 5 points of 15.6% and 7.0%. Gemini ran on a newer model for the whole second window than for most of its baseline.
1 of 552 · 0 of 680 the bare name through the API, in the second window and before it	<pre>node tools/geo-probe/appendix.mjs --paper=entity-records</pre> Rate arm, D03 cells only, control files excluded. The one hit, 2026-10-07, cites the N° 01 paper for the line that the name is widely used.
3,547 of 3,864 · 3,327 of 3,645 the seven other brand questions, second window against 20 August to the first read	<pre>node tools/geo-probe/appendix.mjs --paper=entity-records</pre> The registered negative control: these must not move. 91.8% against 91.3%.
5 of 5 the full name, signed out in the app	<pre>node tools/geo-probe/cold-reads.mjs</pre> Four on 16 September and one on 16 October. Controls, flagged `control: true`, never pooled into a series.
0 of 54 brand answers given without a search that named a fact about us	<pre>node tools/geo-probe/memory-report.mjs --from=2026-09-16 --to=2026-10-15</pre> OpenAI only; Gemini searched on every brand question. A fact the question already contained does not count.
0 of 81 the bare name in the once-a-day runner, three engines, in the second window	<pre>node tools/geo-probe/appendix.mjs --paper=entity-records</pre> Wave arm, D03 only. OpenAI has 23 of the 81 because its account was empty until 22 September.
7,961 requests from named AI crawlers in the second window	<pre>node tools/geo-probe/crawl-report.mjs --from=2026-09-16 --to=2026-10-15</pre> Self-declared user agents; Googlebot, Bingbot, Applebot and Search Console's fetch tool are counted separately. 10.4% of all named-crawler traffic could not be confirmed by the edge.
every figure above the paper against the data	<pre>node tools/research-paper/verify-figures.mjs --paper=entity-records --strict</pre> Re-derives every value on the plates and fails on drift. A value it cannot check reads UNVERIFIED, which also fails.

3 · Every collection day, counted

Answers that cited us, over questions asked, for each of the 29 days. The last two columns are the paper's headline: a zero that holds for 29 consecutive days is a different claim from a zero that is an average.

DATE	NAMES THE COMPANY	TOPIC	BUYING INTENT
2026-09-16	14/16	0/24	0/8
2026-09-17	14/16	0/24	0/8

DATE	NAMES THE COMPANY	TOPIC	BUYING INTENT
2026-09-18	14/16	0/24	0/8
2026-09-19	14/16	0/24	0/8
2026-09-20	14/16	0/24	0/8
2026-09-21	14/16	0/24	0/8
2026-09-22	19/24	0/36	0/12
2026-09-23	20/24	0/36	0/12
2026-09-24	21/24	0/36	0/12
2026-09-25	19/24	0/36	0/12
2026-09-26	19/24	0/36	0/12
2026-09-27	19/24	0/36	0/12
2026-09-28	20/24	0/36	0/12
2026-09-29	19/24	0/36	0/12
2026-10-01	19/24	0/36	0/12
2026-10-02	20/24	0/36	0/12
2026-10-03	19/24	0/36	0/12
2026-10-04	20/24	0/36	0/12
2026-10-05	19/24	0/36	0/12
2026-10-06	20/24	0/36	0/12
2026-10-07	20/24	0/36	0/12
2026-10-08	19/24	0/36	0/12
2026-10-09	20/24	0/36	0/12
2026-10-10	19/24	0/36	0/12
2026-10-11	19/24	0/36	0/12
2026-10-12	19/24	0/36	0/12
2026-10-13	19/24	0/36	0/12
2026-10-14	19/24	0/36	0/12
2026-10-15	19/24	0/36	0/12

4 · The counting decisions behind those numbers

Judgment calls, not arithmetic. Re-derive these figures without knowing them and you will get different answers that are also defensible, which is exactly why they are written down rather than remembered.

1. The second observation window runs 16 September to 15 October. The registration called it “the block” and never dated its start; the date was fixed in the ledger before the paper was drafted. OpenAI has no rows before 22 September, so its figures are the same from either start.
2. The 2026–10–16 recovery run and every file flagged `control` or `excluded\_from\_series` are removed before counting. Nothing was collected on 30 September (no wave or rate file exists for it), so the second window holds 29 wave days and 23 rate-arm days.
3. A citation means an answer links a page on our own domain. The one bare-name hit counts: it links our N° 01 page, even though it does not describe us as one of the Cortexas.
4. The cold reads are scored by hand from the full answer and every follow-up list, never by matching text: an answer can name the state or the word “consulting” and still be about someone else.
5. The second read ran on 16 October, a day after its registered date. It is reported as late, not redated.
6. The 23 namesakes count companies, not web addresses. Cortexa Pty Ltd appeared under two (its own domain and its LinkedIn page) and counts once.
7. The agency and topic questions run by two other studies (845 answers in the second window) are reported only to say they held no citation of us. They are not part of either registered test.
8. The app arm’s earlier bare-name answers (104 cited of 135 in August and September) are not cold reads: that prompt asks all eight brand questions in one message, so “Cortexa” always arrived beside “Cortexa Consulting”.
9. Nothing is pooled across engines or arms. The wave arm, the rate arm and the app are different instruments and each figure names the one it comes from.



WHAT WE SAID BEFORE WE KNEW

Written down first, and empty both times.

The hypothesis, the falsifier and both reading dates were published in our ledger before the intervention. The one API answer that cited a page of ours cited our first paper, and it is reported beside the reads it does not change.

2026-08-25

PREDICTION REGISTERED

the hypothesis, the falsifier, and the 0-of-4 baseline; the second reading booked on 08-27, before any result

2026-09-15

FIRST READING

0 of 4

2026-09-18

SEARCH PREDICTION REGISTERED

0% of non-brand searches name us, ± 5 points

2026-10-16

SECOND READING

0 of 4, a day late; 0 of 654 searches; 1 of 552 API asks, citing N° 01

HOW TO CHECK US

The cold asks are four fixed questions in fresh, signed-out sessions, and success was defined in advance as any page of ours cited or our name in a follow-up list. The search-phrase test reads the same frozen 24 questions the study has asked daily since 7 August.

ABOUT CORTEXA FIELD STUDIES

Cortexa Field Studies is where Cortexa Consulting publishes its own measurement. We run the instrument on our own property first, publish the method before the result, and say plainly when a reading comes back empty.

Intelligence, **applied.**

