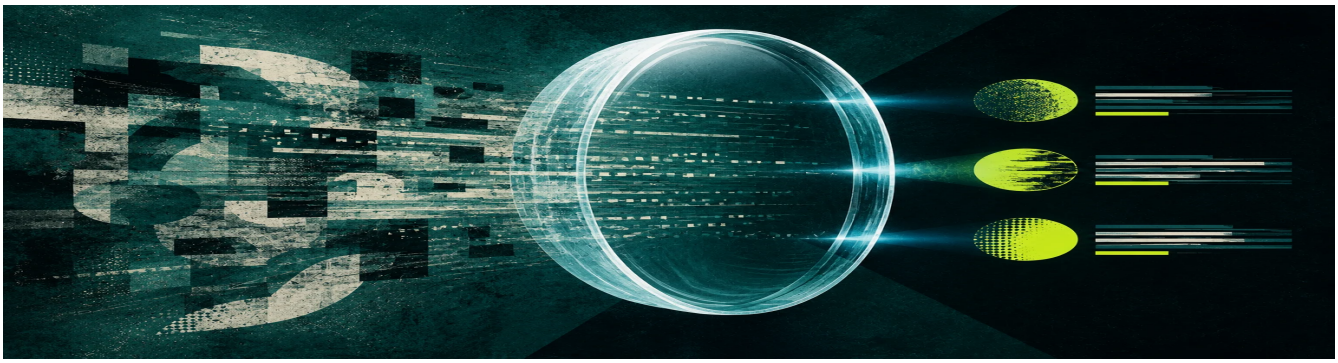
Start with a decision. Keep the evidence.

A practical pilot for marketing teams and business owners who want to understand how their brand appears in AI answers.



What you can do with this pack

Define a research question, organise prompts, capture the answers and turn a finding into a testable action. The PDF explains the method. The companion CSV worksheets open in Excel or Google Sheets.

A pilot, not a market-share claim

A small sample can help you find questions worth investigating. It cannot tell you what every customer sees or establish your share of all AI search. Record the limits before reporting a score.

Write your brief before collecting answers

Decision: What will this research help you decide?
Coverage: Which brand, product, market, language and customer journey?
Comparison: Which alternatives would that customer consider?
Ownership: Who will review the evidence and implement changes?

Start in 01-scope-brief.csv. Keep client material in an appropriate private workspace.

Choose questions that reflect a real journey.

Use customer questions to design the sample. Asking for your brand by name is a different test from seeing whether it is recommended without a prompt.

Journey	Prompt pattern	What it helps explore
Discovery	What options help [audience] solve [problem] in [market]?	Whether the brand enters a relevant consideration set.
Comparison	How do [A] and [B] compare for [constraint]?	How alternatives and trade-offs are represented.
Validation	What should I check before choosing [brand/product]?	Accuracy, suitability and supporting evidence.

Build the register

Give each prompt an ID. Save its exact wording, journey stage, branded or unbranded status, market, language and reason for inclusion. Use real customer evidence where available; label your own hypotheses.

Separate the stable sample from exploration

Keep a defined set unchanged for later comparisons. Put new questions in an exploratory group. Otherwise a new score may simply reflect a different set of prompts.

Choose the collection conditions before testing

Record the platforms, visible search mode, run dates and planned repetitions. Use fresh conversations for independent tests. If you test a multi-turn conversation, save the full context.

Use 02-prompt-register.csv. There is no universal correct prompt count; coverage follows the decision and the variation you need to investigate.

An answer needs a record you can revisit.

One row per observed answer. Store the full answer or a durable capture separately, and put its reference in the evidence log.

Record	What to save
Identity	Run ID, prompt ID, exact input and conversation context.
Conditions	Date/time/timezone, platform, visible model, search mode, account state, language and location where known.
Evidence	Full answer reference and relevant citation URLs. Keep collection errors separate from actual answers.
Observation	Mention, recommendation and owned-source citation: yes, no or unclear. Record competing brands separately.
Accuracy	Accurate, inaccurate, mixed or insufficient evidence, with the factual source used to check the claim.
Review	Reviewer, review date, uncertainties and claims needing further checking.

Do not fill unknowns with assumptions

A platform may not expose the exact model version or location used. Record “unknown” rather than guessing. A failed collection is missing data, not evidence that your brand was absent.

Check the source, not just the citation

Open the cited page and assess whether it supports the statement. A source appearing beside an answer is not proof that it caused the recommendation.

Use 03-evidence-log.csv. Keep identifiers consistent so a finding can point back to its prompt and saved answer.

Keep the measures separate.

A mention, a recommendation and a citation answer different questions. Combining them into one unexplained score can hide what you need to fix.

Measure	Working definition
Mention	The brand is explicitly named in the answer.
Recommendation	The answer presents the brand as an option for the user’s task. A passing reference is not enough.
Owned citation	The answer cites a page on the organisation’s own website.
Accuracy	Check factual claims against reliable information. Use “unclear” or “insufficient evidence” when needed.

Worked calculation: fictional numbers

If 3 of 10 successfully collected answers mention a brand, report “3 of 10 observed answers (30%) in this sample.” If two other attempts failed, disclose them separately. This is not 30% of the AI-search market.

Before presenting a percentage

State the question, numerator, eligible answers, exclusions, dates and test conditions. Show the differences between platforms and journey stages. Repeated runs of one prompt are related observations, not separate customers.

Leave uncertainty visible

Compare the individual answers before interpreting a change. Different wording, platform updates or normal output variation may explain it. A before-and-after difference alone does not demonstrate causation.

Turn the finding into a decision.

A useful action connects an observation, a possible explanation, an owner and a test. Keep those separate so a hypothesis does not become a reported fact.

Step	Fictional software-business example
Observation	A comparison answer cites a competitor’s implementation guide and omits the brand.
Evidence	Save the prompt, full answer, collection conditions and cited guide in the log.
Interpretation	The guide addresses an implementation question. One answer does not establish why the brand was absent.
Investigation	Check whether the pattern recurs. Review whether the brand already provides accessible, accurate implementation information.
Action if confirmed	Ask the content owner and product specialist to improve the relevant page, supported by a verified import example and known limitations.
Validation	Check the published page, record the intervention date and rerun the unchanged sample. Report variation and alternative explanations.

Give the action an owner and a completion test

“Improve content” is not a usable handover. Specify the page, the missing information, who can verify it and how you will check the implementation. Separate the publishing check from later visibility measurement.

Use 04-findings-and-actions.csv. This example is fictional and does not imply a client outcome or a guaranteed route into an AI answer.

Make the next run comparable.

The baseline is a record of the starting conditions. Keep it intact so you can inspect what changed after an intervention.

Before collecting again

1. Freeze the comparison prompts and record any exploratory additions.
2. Record what was implemented, where and when.
3. Verify the implementation before attributing a visibility change.
4. Repeat under documented conditions; note platform changes and unknowns.
5. Show changed answers, unchanged answers and variation alongside totals.
6. Agree what the evidence supports doing next.

Want help designing the full programme?

Bring Rob in to define the sample, investigate the sources, connect technical and content findings, or review your first results. You can use the starter independently and commission support when you need it.

Explore the AI Search Baseline and Action Plan
Discuss your project with Rob

Use and limitations

You may adapt the worksheets for internal work and client delivery. Do not resell or redistribute this starter as a standalone product. No confidential client records are included. A full audit requires research tailored to the business, source investigation, technical evidence and human review.

Current platform guidance

Google Search: AI features and your website
OpenAI: Publishers and developers FAQ