

Make it ask before it builds.

One block of text you paste above any real request. It forces the model to interview you first, show you its plan, and wait for approval before it writes a word.

BY DAVE BROWN

The usual failure with a long AI deliverable is not bad writing. It is a well written document that answers a question you did not ask. You wanted a one page brief for a skeptical CFO and you received two thousand words for a general audience, with an executive summary you did not want and none of the context that would have made it land.

That happens because a model cannot leave a gap open. Every unstated detail gets filled with a default: middle length, neutral tone, generic reader, standard structure. You never see those choices get made. You only see the result, at which point the wrong assumptions are baked into every paragraph and fixing them means starting again.

A model never says “I don’t know enough to start yet.” It just starts.

So make the interview mandatory. Everyone who has briefed an agency, a contractor, or a new hire already knows this move. The good ones ask ten questions before they touch anything. The prompt below is that behavior, written down and made non optional.

6

steps, run in order,
before any drafting
happens

10

questions, delivered at
once, ranked by
consequence

2

places it has to stop and
wait for you

3

minutes of answering
that saves an hour of
rewriting

THE PROMPT

Paste this, then your request at the bottom.

Works in any chat window. Nothing to install, no settings to change. The only part you edit is the last line.

COPY FROM HERE

PLAN MODE – do not produce the deliverable yet.

I'm going to ask you to create a document. Before you write a single word of it, you will run an intake interview. Follow this sequence exactly.

STEP 1 – RESTATE

In two sentences, tell me what you think I'm asking for and what type of document it is (e.g. marketing plan, SOP, financial model, brief, proposal).

STEP 2 – DECLARE YOUR DEFAULT ASSUMPTIONS

List 5–8 assumptions you would make if I gave you nothing else – audience, length, format, tone, time horizon, level of detail, what you'd include and what you'd leave out. Be specific and committal ("I'd assume ~1,500 words, written for an internal ops team, no exec summary"). Vague assumptions are useless to me. Mark any assumption you think is most likely to be wrong.

STEP 3 – ASK ME 10 QUESTIONS

Ask up to 10 questions, numbered, all in one message – do not drip them one at a time. Prioritize by how much the answer would change the output: question 1 should be the one where a wrong guess wrecks the whole document.

Rules for the questions:

- Every question gets a proposed default answer in brackets, so I can reply "3, 5, 7" and let you use your defaults for the rest.
- Ask about things you genuinely cannot infer: internal context, constraints, audience, decisions already made, what's off the table, what "good" looks like, what this replaces or feeds into.
- Do not ask questions you could answer yourself, and do not ask me to do your thinking for you (no "what sections would you like?").
- Include at least one question about what should NOT be in the document.
- Include at least one question about how success will be judged and by whom.
- If a real example, source document, or past version would help, ask for it.

STEP 4 – WAIT

Stop. Do not proceed until I answer. If my answers are thin or contradictory, say so and ask up to 3 follow-ups.

STEP 5 – WRITE THE PLAN

Then produce a plan, not the document:

- The full outline, section by section, with a one-line description and rough word count for each
- Format, tone, and length
- What you're explicitly excluding and why
- Any remaining assumptions you're proceeding on
- Anything you'd need from me (data, names, numbers) that would otherwise be a placeholder – list them
- Open risks: where you think this could come out wrong

Ask me to approve, change, or expand the plan.

STEP 6 – BUILD

Only after I approve, write the document.

My request: [DESCRIBE WHAT YOU WANT HERE]

Your request on the last line can stay short. Two sentences is fine. The whole point is that the model is about to ask you for the rest.

WHAT EACH STEP BUYS YOU

Every step removes one way this goes wrong.

None of the six is decoration. Delete one and a specific failure comes back, usually the one you will not notice until the document is finished.

STEP 01

Restate.

Two sentences that catch the wrong document entirely. This is where you find out it heard "competitive analysis" when you said "competitive positioning," and it costs you ten seconds instead of a thousand words.

STEP 02

Declare the assumptions.

The dangerous assumptions are the invisible ones. Written into a list, they become arguable, and you can strike the wrong one in four words. Asking it to flag the assumption most likely to be wrong is the highest value line in the whole prompt, because it makes the model do your triage.

STEP 03

Ten questions, all at once.

Drip fed questions turn into an interrogation nobody finishes. One batch, ranked so the most consequential is first, is a brief you can answer over coffee. The bracketed defaults are what makes it survive contact with a busy person: reply "3, 5, 7" and it proceeds on its own answers for the rest.

STEP 04

Wait.

Models resolve ambiguity by producing something. Without an explicit stop, many will ask their questions and then helpfully answer them and keep going. One word is carrying real weight here.

STEP 05

The plan, not the document.

An outline is the cheapest thing in the process to fix. You can read one in ninety seconds and see that section four answers the wrong question. You cannot check a finished twelve page draft that fast, which is why weak work often survives simply because it is already complete.

STEP 06

Build.

By this point the model is executing an approved brief rather than guessing at one, and you are reviewing a document you already agreed to the shape of.

THE THREE QUESTIONS DOING THE HEAVY LIFTING

What should **not** be in it. How success gets judged, and **by whom**. And a real past version to work from. Those three are required by the prompt on purpose, because they are the three a model will never think to ask and the three that most often decide whether the output is usable.

RUNNING IT

Two habits, and one honest limit.

Answer sparsely on purpose. You do not owe it ten answers. The bracketed defaults exist so you can answer the two or three that matter and let it commit to the rest. What you get back is a document built on stated assumptions rather than hidden ones, which is the actual win.

Push back on the plan. Step five is a real gate, not a formality. Cutting a section at the outline stage costs a sentence. Cutting it after the draft costs a rewrite, and you will be more reluctant to do it, because by then something has been made.

Skip the whole thing when the work is small. This is overhead, and overhead has to earn its place. Use it when the deliverable is long, when it goes to someone whose opinion matters, or when you would struggle to check the result yourself. For a quick draft you are going to rewrite anyway, just ask for the draft.

WHAT WENT WRONG LAST TIME

THE STEP THAT PREVENTS IT

It wrote the wrong kind of document

Step 1, and you saw it in two sentences

Right topic, wrong reader

Step 2, where audience becomes a stated choice

Generic because it lacked your context

Step 3, the questions it cannot answer itself

Included the thing you specifically didn't want

Step 3, the required question about exclusions

Structurally wrong, so the fix was a rewrite

Step 5, caught at outline cost

Full of confident placeholders

Step 5, which lists what it needs from you first

The model isn't missing intelligence. It's missing your context, and it will never ask for it unless you make it.

Left alone, it fills every gap with a default and shows you the finished result. This prompt moves the gaps to the front, where they cost a sentence each, and puts two gates between the request and the draft. Three minutes of answering, one look at an outline, and the document you get back is the one you meant to ask for.

More guides like this one, for people trying to use AI without embarrassing themselves.

Weekly, plain-language breakdowns on Instagram.

[@davecto](#)