

The "Grill Me" Trick for AI

Here's a simple habit that makes any AI chat (Claude, Gemini, whatever you're using) dramatically more useful: instead of asking it to *help* you, ask it to *interrogate* you.

Why bother

LLMs are agreeable by default. If you bring them a half-formed plan, they'll happily polish it and tell you it's great — which feels nice and teaches you nothing. "Grill Me" flips that. You turn the AI into a skeptical interviewer whose only job is to poke holes in your thinking *before* you've sunk time, money, or code into it.

It's most valuable right before you commit to something: starting a project, writing up a plan or spec, making a decision, or sending something that matters. Ten minutes of being grilled tends to surface the assumptions you didn't know you were making.

How to do it

Paste this at the start of a session, filling in the blank:

```
I want you to grill me on [the thing I'm about to do or decide].
```

```
Act as a sharp, skeptical interviewer whose job is to pressure-test my thinking before I commit to it. Ask me ONE question at a time. Each question should probe an assumption, expose a gap, or force me to be more specific. Wait for my answer before asking the next one. Don't be agreeable — push on weak reasoning, vague goals, and anything I'm hand-waving past.
```

```
Keep a running record as we go. After each answer I give, and before you ask the next question, append a short entry to a markdown file called grillme-log.md — the question, my answer, and any decision we reached. If you can't write files in this tool, print that entry as a markdown code block each turn instead, so I can save it myself.
```

```
When you've covered the important angles, give me a summary of what's solid, what's still shaky, and what I haven't figured out yet — and add that summary to the same record.
```

That's the whole thing, and it works in any tool — though *where* the record lives depends on which one.

Why the running log: in a long session the AI quietly forgets the early part of the conversation to make room — so if you wait until the end to ask for a summary, your first and often best decisions can already be gone. Writing each one down as you go means nothing important evaporates before you're finished. In Claude Code it'll save the file for you automatically; in a chat like Gemini it'll show you each entry to keep.

Steering it while it goes

You're in charge of the conversation, so redirect freely:

- "Go deeper on that one."
- "Move on, that's settled."
- "You're being too easy on me — push harder."
- "That's out of scope for now."
- "Focus on the [budget / timeline / technical] side."

Knowing when you're done

There's no buzzer — it's a feel. You're done when the questions start landing softly: you can answer them crisply without hesitation, or they're getting repetitive and trivial. The AI's closing summary helps too — when the "still shaky" list has shrunk to a few things you've *consciously decided* to live with, you've gotten what you came for. The goal isn't zero unknowns; it's understanding your own plan better than when you started.

One more thing

Some tools (Claude Code, for example) let you save a pattern like this as a reusable "skill" or saved command, so it's always on tap instead of something you paste in. Nice to have eventually, but honestly overkill here — this prompt is short enough that pasting it does the same job everywhere.