ERICFONDREN.COM

Second-opinion audit

A fresh, skeptical review of your plan, code or decision, checked before you see it. By Eric Fondren

What it does

You point your AI at a piece of work: a plan, a document, a code change, or a decision you are weighing. It hands that work to a fresh reviewer that has seen none of your conversation and has no stake in the work being right. The reviewer's only job is to find what is wrong and show the evidence.

Then your AI checks each serious finding against the source before showing you anything. You get a short report that leads with how many findings held up, sorted into three groups:

- **Confirmed.** The evidence backs it up.
- **Plausible, untested.** It might be real, and here is what would settle it.
- **Dropped.** The reviewer got it wrong, with the reason.

Nothing changes in your files unless you ask.

An example. Jordan runs a small dog-grooming shop. She has a one-page plan to move bookings from a shared spreadsheet to an online booking tool next week. She types "get a second opinion on this plan". A few minutes later:

> 6 findings, 4 held up.
>
> **Blocker, confirmed.** The plan turns off the spreadsheet the same day the booking tool goes live. If the import fails, there is no list of next week's appointments. Fix: run both side by side for one week.
>
> **Major, confirmed.** Regular customers book by text today. The plan never says who tells them, or how.
>
> **Major, plausible, untested.** The tool may not handle a single appointment that books two groomers at once. Check its settings before switching.

(two more confirmed findings, then:)

Dropped: "No one owns the cutover." The plan names Sam for it in step 3.

What you need

- An AI assistant: **Claude Code**, **Codex**, **Claude.ai** or **ChatGPT**. Any plan you already pay for works. No extra account.
- For reviews of code: a project tracked in git. Your AI can check this for you.
- Optional: if you also have another company's AI command-line tool (such as Codex or Gemini), the same review can go to a different AI model for a second view.

In Claude Code and Codex, the reviewer runs on its own inside your assistant. In Claude.ai and ChatGPT chat, your AI gives you a message to paste into a new chat, and you paste the answer back. Both get the same result. The chat route takes one extra copy and paste.

How to start

1. Download the zip from the tool card.
2. Give it to your assistant and say: **"Set up this skill using its README."** It installs the skill and runs a short test.
3. Start a review:
    - Claude Code: type **/second-opinion**
    - Codex: type **$second-opinion**
    - Any assistant: say **"get a second opinion on this"** and point at the plan, file, change or decision.

If your chat assistant cannot install skills, paste the all-in-one file into a chat first, then ask for a second opinion.

What to expect

It asks you nothing about yourself and keeps no profile. It may ask two things:

- **Which work** to review, if that is unclear.
- **What is at stake**, if it cannot tell. For example: "this touches checkout, so a bug loses orders" or "we will be locked into this vendor for two years." This tells the reviewer which problems matter most.

A review takes a few minutes for a short plan and up to fifteen or more for a large code change. Your AI tells you when it starts and when it is checking the results.

It reviews three kinds of work, each with its own checklist:

- **Plans and documents.** Unstated assumptions, missing steps, steps with no owner, what happens if one piece runs late, how to undo it.
- **Code changes.** Two things happening at once, crashes halfway through, removed code that something still uses, checks that can be skipped, cost and speed at real volume.
- **Decisions.** What has to be true for this to be right, the cost of being wrong, options you have not considered, and who else is affected.

At the end it offers to fold the confirmed findings into your plan or fix the code. It waits for your yes.

Tips and limits

- **Write it down first.** The reviewer only sees what it is given. A plan that lives in your head gets a thinner review. If you talked a plan through with your AI, it writes that out word for word for the reviewer.
- **Say what is at stake.** One sentence about what goes wrong if this fails sharpens every finding.
- **Findings are leads.** Your AI checks the serious ones, but "plausible, untested" means exactly that. Settle those before you act on them.
- **For decisions, it only knows what you tell it.** If a finding depends on a fact you did not share, it says so and names the question to answer.
- **Same model, fresh eyes.** The reviewer usually runs on the same AI model as your assistant. Its independence comes from the clean start. For a different model's view, use the optional command-line route above.
- **In the chat route, use a brand-new chat.** A temporary or incognito chat is best, outside any project folder, so the reviewer cannot see your earlier conversation.
- **If it stalls,** ask your AI for the review prompt and run it yourself in a new chat. Paste the answer back and ask it to triage.
- **Ask for a re-audit** after you revise. The second round focuses on what changed and on anything still serious.