

The Board: stop your AI just agreeing with you

Ask one AI a question and you get one answer, almost always the one you were hoping for. The Board runs your question past five advisors who each attack it from a different angle, makes them review each other, then a Chair hands you a straight verdict. Here is the whole thing, plus the Claude skill.

Your AI has a yes-man problem

You ask your AI if your idea is good. It says yes. You feel great. Six months later the idea falls over on a flaw anyone with a different perspective would have spotted in five seconds. This is not a hunch. A Stanford study published in *Science* in March 2026 put numbers on it.

Researchers Myra Cheng and Dan Jurafsky tested 11 leading models (GPT-4o, GPT-5, Claude, Gemini, DeepSeek and others) across nearly 12,000 prompts and 2,400 people. What they found:

49% more agreeable than a human

The models endorsed the user far more readily than a person would.

51% sided with the user even when the crowd said they were wrong

Using Reddit's "Am I The A-hole" posts where humans had already voted someone wrong, the models still took the poster's side over half the time.

It gets worse. A single chat with a sycophantic AI left people less likely to apologise, less likely to admit fault, and more sure they were right. And they were more likely to keep using the agreeable version. The thing that makes your decisions worse is also the thing that feels best to use, and when asked, people could not tell the flattering version from the honest one. You do not notice it happening.

The fix: a board, not a chatbot

One AI gives you one perspective, and that perspective was built to agree with you. The Board fixes that. It runs your question through five independent advisors, each thinking from a completely different angle. Then they review each other's work. Then a Chair pulls it together into a verdict that tells you where they agree, where they clash, and what to actually do.

This borrows from Andrej Karpathy's "LLM Council" idea: send a question to several models and have them peer-review each other anonymously. The Board does the same thing inside Claude using five different thinking lenses.

The five advisors

These are not job titles or personas. They are thinking styles that naturally pull against each other.

1. **The Skeptic.** Hunts for what breaks. Assumes there is a flaw and goes looking for it. Not negativity for its own sake, the friend who stops you signing a bad deal by asking the question you have been dodging.

2. **The Reframer.** Ignores how you asked it. Asks what you are actually trying to solve, strips the assumptions, rebuilds the problem from scratch. Will tell you when you are solving the wrong thing entirely.
3. **The Visionary.** Looks for the upside everyone else is ignoring. What if this is bigger than you think? What adjacent move is hiding here? Does not touch risk, that is the Skeptic's job.
4. **The Newcomer.** Knows nothing about you or your field. Reacts only to what is in front of them. Catches the things obvious to you and baffling to everyone else. The most underrated voice in the room.
5. **The Operator.** Cares about one thing: can this be done, and what is the first move? No theory. Looks at everything through "what do you do Monday morning?"

Three natural tensions fall out of this. Skeptic versus Visionary (downside versus upside). Reframer versus Operator (rethink it versus just do it). The Newcomer sits in the middle keeping everyone honest with fresh eyes.

How a session runs

You start your message with "ask the board" and the rest happens on its own.

- **Phase 1, the five answer.** Each advisor writes 150 to 200 words from their angle. No hedging, no balance. The Skeptic goes hard on flaws, the Visionary goes hard on upside. No middle ground.
- **Phase 2, peer review.** This is what makes it more than asking five times. It looks across all five and asks: who made the strongest case, who has the biggest blind spot, and what did all five miss? The review is where the real insight shows up.
- **Phase 3, the Chair's verdict.** A final synthesis in five parts: where the board agrees, where it clashes, blind spots caught in review, the call (one clear recommendation, not "it depends"), and one thing to do first.

The Chair can overrule the majority. If four advisors say "do it" but the lone dissenter has the better argument, the Chair sides with the dissenter and explains why. This is not a vote. It is about the best answer.

The prompt

This is the whole system. Paste it into a Claude Project's instructions (the recommended setup is below).

The Board

When I say "ask the board", "run it past the board", or "board this" before a question, run my question past five independent advisors. Each one thinks from a different angle. After they answer, run a peer review, then a Chair gives the final verdict. Do not narrate the machinery, just run it.

The five advisors

Write a response from each of these five (150 to 200 words each). Do not hedge and do not try to be balanced. Each one leans all the way into their angle.

1. THE SKEPTIC

Hunts for what breaks. Assumes there is a flaw and goes looking for it. Not negativity for its own sake. The friend who stops you signing a bad deal by asking the question you have been dodging.

2. THE REFRAMER

Ignores how I asked the question. Asks what I am actually trying to solve, strips out the assumptions, rebuilds the problem from scratch. Will tell me when I am solving the wrong thing entirely.

3. THE VISIONARY

Looks for the upside everyone else is ignoring. What if this is bigger than I think? What adjacent move is hiding here? Does not touch risk, that is the Skeptic's job.

4. THE NEWCOMER

Knows nothing about me, my work, or my history. Reacts only to what is in front of them. Catches the things obvious to me and confusing to everyone else.

5. THE OPERATOR

Cares about one thing: can this be done, and what is the first move? No theory, no big-picture talk. Looks at everything through "what do I do Monday morning?" If there is no clear first step, says so.

Peer review

After all five answer, review them. Decide:

1. Which advisor made the strongest case, and why.
2. Which one has the biggest blind spot.
3. What all five missed and should have caught.

The Chair's verdict

Pull it together into a verdict with these exact sections:

WHERE THE BOARD AGREES – points more than one advisor reached on their own. High-confidence signals.

WHERE THE BOARD CLASHES – the real disagreements. Both sides, no smoothing over.

BLIND SPOTS CAUGHT – what came out of the peer review that no single advisor caught.

THE CALL – one clear recommendation. Not "it depends". A real answer.

DO THIS FIRST – one concrete next step. Not a list. One thing.

Rules

- Never agree with me just to keep me happy.
- Each advisor leans all the way into their angle.
- The Chair can overrule the majority when the lone dissenter has the better argument. This is not a vote.
- If my question is vague, ask me one sharp question before convening the board.
- Do not convene the board for trivial questions with one right answer.

Three ways to set it up in Claude

1. Claude Project (recommended)

A Project gives you a dedicated space whose instructions apply to every chat inside it. Project instructions hold around 8,000 characters, so the full prompt fits with room to spare.

1. In Claude, open the sidebar and click *Projects*, then *New Project*. Name it "The Board" or "Decisions".
2. Open the Project's instructions and paste the full prompt above.
3. Start any chat inside that Project with "ask the board" followed by your question.

2. Custom instructions (everywhere, condensed)

If you want the Board in every chat, not just one Project, add a shorter version to your account custom instructions. In Claude, open *Settings*, go to your profile, and find *custom instructions*. That field caps at about 1,500 characters, so use this trimmed version. It works, but the Project version is more detailed and gives better results.

THE BOARD: CONDENSED (CUSTOM INSTRUCTIONS)

When I say "ask the board" or "board this", run my question past 5 advisors before answering:

1. The Skeptic: find what breaks, what I am dodging. Assume a flaw exists.
2. The Reframer: ignore my framing. What am I actually trying to solve?
3. The Visionary: find the upside everyone missed. What if this is bigger?
4. The Newcomer: you know nothing about me. What is obvious to a stranger?
5. The Operator: skip theory. What do I do Monday morning? If there is no first step, say so.

Write about 150 words each. Be direct, no hedging, lean all the way into each angle.

Then peer review: who is strongest, who has the biggest blind spot, what did all 5 miss?

Then a Chair's verdict:

- Where the board agrees (high-confidence signals)
- Where it clashes (the real disagreements)
- Blind spots caught in review
- The call (one clear recommendation, not "it depends")
- Do this first (one concrete next step)

Never agree just to keep me happy. The Chair can overrule the majority if the reasoning is stronger.

3. Install it as a Claude skill (best)

If you use Cowork (Claude's desktop app), the Board runs as a real skill instead of pasted text. Each advisor runs as its own sub-agent, so they genuinely cannot see each other until the peer review, which catches more blind spots. Download the skill file and, in Cowork, go to *Settings*, then *Capabilities*, then *Skills*, and upload it. After that, "ask the board" works in any conversation.

[Download the Board skill ↓](#)

See it in action

A real question where being wrong is expensive: *"Ask the board: I want to raise my course price from \$97 to \$297. My audience is non-technical founders. Smart move?"*

- **The Skeptic:** at \$297 you are now competing with free content and you raise refund risk with a non-technical crowd. The people who would pay \$297 may already be past beginner.
- **The Reframer:** what is the price even for? Revenue, authority, or feeding a higher-ticket offer? The right number changes completely depending on which.
- **The Visionary:** beginner founders are underserved while everyone teaches advanced stuff. Done right, \$297 might be too low, not too high.
- **The Newcomer:** I would not know if this course is for me from the title. Sell the outcome, not the topic.
- **The Operator:** do not rebuild anything yet. Run one live session at the new price to 30 people first. If they do not buy, the course will not either.

The Chair's call: the price is not the real question, the positioning is. Test \$297 with a single live workshop before you touch the course. **Do this first:** sell 30 seats to a live session at the new price and watch the conversion before committing.

When to run it

You do not need the Board for everything. Save it for decisions where being wrong costs you time, money, or reputation.

- **Good:** "Should I raise my prices?" "Which of these 3 positioning angles is strongest?" "I'm thinking of pivoting from X to Y, am I crazy?" "Here's my landing page, what's weak?"
- **Skip it:** anything with one right answer. "What's the capital of France?" "Write me a tweet." "Summarise this."

The rule of thumb: if you catch yourself hoping the AI agrees with you, that is exactly when to run it past the board. The point is the truth, not the validation.

That is the build. If it was useful, the newsletter is where I send the next one before it hits YouTube.

Built for curious people, not experts.

© 2026 Vasco Yaps. All rights reserved.