

# The Advisor Interview

A best-practices guide for capturing how a person thinks, decides, and pushes back — so their judgment can keep advising you.

---

**Format:** 1 conversation, ~90–120 min    **Record:** audio + verbatim transcript    **Draft** v1 · for review

---

This is a single, semi-structured conversation. Its goal is not a résumé or a list of opinions — it is to capture the **reasoning underneath** them: how this person weighs a decision, what they reliably doubt, what they actually know cold, and how they say it. The transcript is the raw material; the questions below are a map, not a script.

## For the participant

If someone handed you this, here is the honest version of what it is and how it will be used.

I want to keep learning from how you think — even on the days I can't borrow an hour of your time. With your permission, I'd like to record a conversation and use it to build a **private, AI-assisted "sounding board"** that reflects your perspective, so when I'm working through a hard decision I can pressure-test it against a model of your judgment.

### A few promises:

- It is for **my own thinking only**. It will never be used to impersonate you, or to speak as you to anyone else.
- You can **skip any question**, go off the record at any point, or redact anything afterward.
- I'll share back what I build if you're curious, and **delete the recording** whenever you ask.
- What's most useful to me is your *reasoning* — not polished conclusions. Half-formed is fine. Disagreeing with the question is better.

---

PARTICIPANT NAME

---

CONSENT TO RECORD · DATE

## How to run it Interviewer notes

The quality of the model is set here, not later. A few rules that consistently separate a rich transcript from a thin one:

- **Listen ~80% of the time.** Your job is the follow-up, not the speech. Silence earns better answers than the next question does.
- **Don't lead.** Never signal the answer you expect — it contaminates the conversation now and the model later. Ask, then wait.
- **Chase stories, not abstractions.** "Tell me about a time..." beats "What do you value?" When you hear a principle, ask for the example; when you hear an example, ask for the principle.
- **Follow the energy.** Go off-script when they light up. The map is there to return to, not to obey.
- **Capture exact wording.** Their phrases and metaphors *are* their voice — keep them verbatim, don't paraphrase.
- **Mark where the knowledge stops.** Note what they defer on. That boundary is what later keeps the model honest instead of confidently wrong.

### WHY THESE SIX PARTS

Parts 1–2 establish story and decision-making; Part 3 captures how they *doubt* (the most useful and most easily lost signal); Part 4 anchors real expertise so the model can be corrected against facts; Part 5 captures voice; Part 6 sets aside held-out answers you'll use to *test* the result. Aim for the times shown, but trade freely between them.

**PART 1****Story & formation** ~15 min**INTERVIEWER**

Open wide and let them run. Your only moves here are "what happened next?" and "why did that matter to you?"

- 1 Walk me through your path. Where did it start, and what are the two or three turning points that made you who you are in your work?

---

- 2 Tell me about an early experience that shaped how you see your field — one that still sits with you.

---

- 3 Who shaped how you think — mentors, rivals, authors? What did each of them teach you?

---

- 4 What's a belief you held strongly early on that you later changed your mind about?  
What changed it — an event, a person, evidence?

---

- 5 Among people who know your work well, what are you known for?

**INTERVIEWER**

Push for the rule behind the story. After each example: "So what's the general principle there?" Capture the *order* in which they weigh things.

- 1 When you face a hard call in your work, what are you optimizing for — and what are you willing to give up to get it?

---

- 2 Walk me through a recent difficult decision, start to finish. What did you weigh, and in what order?

---

- 3 What's a rule of thumb you trust that many others in your field don't?

---

- 4 When a new idea or proposal lands on your desk, what are the first three things you check?

---

- 5 What's the tradeoff you keep returning to?  
Speed vs. rigor, cost vs. quality, consensus vs. conviction...

---

- 6 What would you never do, even under real pressure?

**PART 3****Convictions & doubts** ~20 min**INTERVIEWER · THE CORE**

This is the most valuable section and the easiest to lose. Capture *how* they object, not just what they reject — the exact question they'd ask, the tell that tips them off.

- 1 What's a widely-held view in your field that you think is wrong, or far too simple?

---

- 2 When someone brings you a plan, what's the first weakness you tend to spot?

---

- 3 What's the question you always ask that other people forget to?

---

- 4 What makes you lose confidence in an idea fast?  
What's an instant red flag?

---

- 5 Tell me about a time you were the lone skeptic and turned out right. What had you seen that others hadn't?

---

- 6 And a time your doubt was wrong — what did that teach you about when to trust an idea anyway?

---

- 7 If you were reviewing my work and wanted to save me from myself, what would you be watching for?

**INTERVIEWER · THE ANCHOR**

Get specifics: numbers, named sources, real examples. When they assert something, ask "where would I go to verify that?" The aim is to separate what they *know* from what they *suspect*.

- 1 Explain how your field really works to a smart outsider — the model in your head that the textbooks don't quite capture.
- 2 What do even experienced people most often get wrong about your domain?
- 3 What are the live debates or open questions right now? Where do you stand, and why?
- 4 What numbers, benchmarks, or facts do you carry around that anchor your judgment?
- 5 Whose work do you trust? What do you read, and what do you deliberately ignore?
- 6 What's changed in the last few years that outsiders haven't caught up to yet?
- 7 Where does your own expertise run out — what would you hand to someone else, and to whom?

This boundary is gold: it's what stops a model from bluffing in your voice.

**PART 5****Voice** ~10 min**INTERVIEWER**

You'll capture most of this just by listening across the whole session. These questions surface the rest.

- 1 How would your closest colleagues describe the way you communicate?
- 2 What are phrases or expressions you catch yourself using a lot?
- 3 What kind of language or framing turns you off the moment you hear it?
- 4 When you disagree with someone you genuinely respect, how do you say it?

**PART 6****Calibration** ~10 min**INTERVIEWER · DO NOT SKIP, DO NOT REUSE**

Record these answers **verbatim and set them aside** — keep them out of the persona's notes. Later, put the same questions to the finished model *cold* and compare. This held-out set is your fidelity test: it's the difference between "it sounds like them" and "it predicts them." Prepare 2–3 short, realistic scenarios from your own domain in advance.

- 1 Here's a short, realistic scenario from your world — what would you advise, and what's the first thing you'd want to know?  
[Prepare 2–3 of these.]
- 2 I'll describe a plan in two sentences. Tell me the first thing you'd challenge.
- 3 Rank these three priorities the way *you* would, and tell me why the bottom one is last.  
[Insert three real, competing priorities.]
- 4 A quick gut check:  
[a pick-one question in their domain with a knowable, recordable answer].

## Closing

- 1 What didn't I ask that I should have?
  - 2 Who else thinks like you — or usefully unlike you — that I should learn from?
- 

Then close the loop: thank them, restate how the recording will and won't be used, and confirm they're still comfortable. Within a day, send a one-line thanks and the offer to redact anything on reflection.

---

### **The Advisor Interview** · best-practices interview protocol · draft v1.

Method grounded in 2024–2026 research on building faithful individual personas: a single ~2-hour semi-structured life interview (life story + views on contested issues, with adaptive follow-ups) is the validated minimum-viable input, materially out-performing demographic or stereotype priors (Park et al., *Generative Agent Simulations of 1,000 People*, Stanford/DeepMind). Fidelity is judged on predicted behavior, not stated traits — hence the held-out calibration in Part 6. Participant consent reflects current digital-duplicate ethics guidance.