— A FIELD GUIDE

The top 10 prompting hacks for *leaders, boards & investors.*

How senior decision-makers get sharper, faster, more reliable answers out of AI — without learning to code, write XML, or read a single API doc.

CURATED BY

Raffaela Rein

FROM

The AI Leadership Edge

— WHY THIS GUIDE EXISTS

Most leaders are still *under-using AI* — not because the model is wrong, but because the prompt is.

Anthropic's Applied AI team trains engineers on a 10-layer framework for production prompts. The framework is excellent. It's also written for engineers shipping code into Swedish car-insurance pipelines.

This guide adapts that same framework for a different reader: the board director reviewing a 200-page board pack on a Sunday night. The general partner triaging twelve portfolio updates before a Monday IC meeting. The CEO who needs to brief their leadership team on a regulatory change by lunch.

None of you will write a line of XML. None of you should have to. But the underlying logic — what the model needs to give you a usable answer — applies just as much to your work as it does to anyone else's. Possibly more, because the stakes are higher and the time is shorter.

Each hack below is one page. A before-prompt, an after-prompt, why it matters, and one concrete scenario for someone in your seat. Read it in 20 minutes. Apply it for the rest of your career.

WHO THIS IS FOR

BUSINESS LEADERS

CEOs, executives, and operators using AI to compress strategy work, brief teams, and prepare for high-stakes conversations.

BOARD DIRECTORS

Non-executive directors and chairs using AI to read board packs faster, spot pattern signals, and ask sharper questions.

GPS & INVESTORS

Partners using AI to triage deal flow, monitor portfolios, and pressure-test memos before they hit the IC.

— THE FRAMEWORK AT A GLANCE

Ten layers. *One good answer.*

Think of a prompt as a brief you'd give a sharp new hire on day one. Skip a layer and they fill the gap with a guess. The same is true for the model.

THE BRIEF — set the stage

1. Set the scene
2. Tell it how to handle doubt
3. Give it the source material

THE METHOD — guide the thinking

4. Sequence the steps
5. Show, don't just tell
6. Restate right before you ask

THE OUTPUT — shape the answer

7. Make it think out loud
8. Demand the shape you want

THE VERIFY — pressure-test it

9. Stress-test the answer
10. Iterate, don't restart

Most lazy prompts use 1–2 layers. Most production-quality prompts use 8–9.

HOW TO USE THIS GUIDE

You don't need all ten on every prompt. A one-line factual question doesn't need a framework. But anything that matters — a briefing, a memo draft, a board summary, a deal screen — should use most of them. The hacks are presented in the order you'd apply them when writing a prompt from scratch.

01 · THE BRIEF

Set the scene.

Tell the AI three things before anything else: **who it is**, **who you are**, and **what situation you're operating in**. Without this, the model fills the gap with the most statistically average interpretation it can find — which is almost never yours.

LAZY VERSION

```
Summarise this board pack.
```

SHARP VERSION

```
You are a senior board advisor with 15 years of experience on Series B+ SaaS boards. I am a non-executive director sitting on the audit committee. I'm reviewing this 80-page board pack before tomorrow's meeting and need to walk in able to ask three sharp questions of the CFO.
```

> *Why it matters: the same model can play "intern" or "former Big Four audit partner" depending entirely on how you frame it. The framing costs you one sentence and changes the entire shape of the answer.*

SCENARIO · GP TRIAGING A DEAL MEMO

"You are a partner at a top-decile Series A fund focused on B2B SaaS. I'm preparing for an IC meeting tomorrow morning on a deal where I have reservations. Read the memo below and tell me the three weakest claims a sceptical partner would push on."

Tell it how to handle doubt.

The default behaviour under uncertainty is to produce a fluent answer that papers over the gap. For high-stakes work that's the failure mode you most need to prevent. Tell it explicitly what to do when it doesn't know.

LAZY VERSION

```
What was the company's revenue growth last quarter?
```

SHARP VERSION

```
Stay strictly factual. If a number isn't explicitly stated in the document I gave you, say "not stated in the source" rather than estimating. If a claim is ambiguous, flag it as "unclear — needs verification" and quote the source sentence. I'd rather have an incomplete answer than a confidently wrong one.
```

> *Why it matters: this is the single highest-leverage instruction for reducing hallucination on substantive work. It costs you nothing and changes everything.*

SCENARIO · BOARD DIRECTOR READING A CFO UPDATE

"Extract every quantitative claim in this CFO report. For each, state whether it's (a) supported by a number in the appendix, (b) directional only, or (c) unverifiable from the document. Do not estimate or fill gaps. List anything you'd flag for the audit committee."

03 · THE BRIEF

Give it the source material.

Models can't see what you don't share. Paste in the actual document, the actual numbers, the actual transcript. Don't paraphrase. Don't summarise. The model is better at reading 50 pages of dense PDF than you are — let it.

LAZY VERSION

```
The CEO said something concerning about
churn in the last all-hands. What should I
do?
```

SHARP VERSION

```
Below is the full transcript of the CEO's
Q3 all-hands (8,400 words). I'm a Series B
investor with a board observer seat. Read
it end to end and identify any statement
about churn, retention, or net revenue
retention. For each one, quote the exact
sentence, note the surrounding context,
and rate whether it raises a flag (low/
medium/high) for an investor.

[PASTE FULL TRANSCRIPT HERE]
```

> *Why it matters: leaders systematically under-share context because they're used to briefing humans who get tired. Models don't get tired. Give them more than you think they need.*

SCENARIO · CEO PREPARING FOR AN ANALYST CALL

Paste in the last three earnings transcripts, the most recent investor letter, and the analyst's prior published note. Then ask: "Identify the three questions this analyst is most likely to ask, based on patterns in their prior coverage and themes they've returned to repeatedly."

04 · THE METHOD

Sequence the steps.

It's not enough to tell the model *what* to do. Tell it the *order*. Without sequencing, the model jumps to the visually salient or most-recent piece of context and reasons backward from there. That's how you get wrong answers on hard cases.

LAZY VERSION

```
Read this memo and tell me if we should
invest.
```

SHARP VERSION

```
Work through this in exact order. Don't
skip steps and don't change the order.

Step 1: List every quantitative claim made
in the memo, with the source page.
Step 2: List every assumption the deal
thesis relies on. Mark each as "stated" or
"implicit."
Step 3: Identify the three biggest risks
only after completing steps 1 and 2.
Step 4: Give your final read on whether
the thesis holds up, grounded in specific
items from steps 1–3.
```

> *Why it matters: this mirrors how the best partners or board members actually think. They don't lead with a verdict — they build to one. The model will do the same if you tell it to.*

SCENARIO · BOARD DIRECTOR PREPARING FOR AN EXECUTIVE REVIEW

"Step 1: Extract the three KPIs the CEO highlighted this quarter. Step 2: Compare each to the trajectory promised in the prior two quarterly reports. Step 3: Only after that, flag any narrative shifts in how the CEO is framing performance. Step 4: Suggest the three questions I should ask in the executive session."

Show, don't just tell.

For any task with judgement involved — what counts as a red flag, what tone you want, what "good" looks like — give the model one or two examples. One concrete example outperforms five paragraphs of abstract instruction.

LAZY VERSION

```
Rewrite this in a more executive tone.
```

SHARP VERSION

```
Rewrite this paragraph in the same voice
as the example below.

Example of the voice I want:
"The quarter missed plan on revenue but
beat on margin. The miss was concentrated
in two enterprise deals that slipped, not
a broader pipeline issue. We expect both
to close in Q1."

Notice: short sentences. Quantitative
claims first, narrative second. No hedging
language. Now apply that voice to my draft
below.
```

> *Why it matters: "executive tone," "investor-grade analysis," "board-ready summary" — these mean different things to different people. The model can match a voice exactly if you show it one. It cannot guess.*

SCENARIO · GP DRAFTING AN LP UPDATE

Paste a paragraph from your fund's best-received prior LP letter. Then: "Write the Q3 portfolio commentary in this exact voice. Same sentence length, same balance of confidence and acknowledgement of risk, same use of specific numbers."

06 · THE METHOD

Restate right before you ask.

In a long prompt with context, examples, and instructions, the actual task can get buried. Right before the data — or right before you hit send on a complex prompt — restate what you want in one line. Yes, it feels redundant. That's the point.

LAZY VERSION

```
[3,000-word prompt with role, context,
documents, and rules]
What do you think?
```

SHARP VERSION

```
[3,000-word prompt with role, context,
documents, and rules]

Reminder: your job is to produce the three
sharpest questions I should ask the CFO
tomorrow, grounded in specific items from
the documents above. Not a summary. Not a
strategy memo. Three questions, with the
source page reference for each.
```

> *Why it matters: the longer the prompt, the more important this line becomes. It's also where you re-emphasise the "no guessing" instruction from Hack 2, because the model's attention drifts over long inputs just like yours does.*

SCENARIO · ANYONE RUNNING A COMPLEX PROMPT

After a long prompt with multiple documents and rules: "One more time — I need exactly three items, in priority order, each one tied to a specific source quote. If you can't find three, give me two. Don't pad."

07 · THE OUTPUT

Make it think out loud.

Ask the model to walk through its reasoning *before* it gives you the answer. This is the single biggest accuracy lever available to non-technical users. It costs an extra few seconds and it works.

LAZY VERSION

```
Which of these three candidates should we hire?
```

SHARP VERSION

```
Before giving me a recommendation, think through this step by step. For each candidate, lay out: (1) what the role needs, (2) what their CV evidences, (3) what's missing or unclear, (4) the strongest case for them and the strongest case against. Only after that, give me a recommendation with reasoning grounded in your analysis above.
```

> *Why it matters: the model produces dramatically better answers when it shows its work. As a bonus, you can read the reasoning and catch where it went off — which lets you correct the prompt rather than the answer.*

SCENARIO · BOARD DIRECTOR EVALUATING A CEO SUCCESSION SHORTLIST

"Walk through each candidate against the four criteria the board agreed on (scale experience, capital efficiency, M&A track record, culture fit). For each criterion, rate the evidence as strong / mixed / thin and cite specific items from their bio. Only then rank the three candidates with reasoning."

Demand the shape you want.

If you don't specify a format, you get prose. Prose is hard to scan, hard to action, and hard to share. Tell the model exactly how to package the answer — a table, three bullets, an executive summary in 80 words, a one-pager with these four sections.

LAZY VERSION

```
Compare these three vendors and tell me
which is best.
```

SHARP VERSION

```
Produce your answer in this exact
structure:

1. A comparison table with columns:
Vendor / Price / Implementation time /
Reference customers / Top risk.
2. A 60-word executive summary directly
below the table, no longer.
3. A "what I'd ask next" section with
exactly three follow-up questions.

No introduction, no caveats, no closing
paragraph. Just those three sections.
```

> *Why it matters: the format is the deliverable when you're forwarding it to others. A board-ready one-pager and a brainstorm paragraph contain the same information but produce wildly different outcomes.*

SCENARIO · CEO BRIEFING THE LEADERSHIP TEAM

"Output as a Slack-ready post: one headline sentence (under 12 words), three bullets with the key implications, one line stating what we're deciding and by when. Total under 100 words."

09 · THE VERIFY

Stress-test the answer.

Once you have a draft answer, don't accept it. Make the model argue against itself. This is what the best directors and partners do instinctively in a real conversation — and you can ask the model to do it in 20 seconds.

LAZY VERSION

```
Great. What else should I know?
```

SHARP VERSION

```
Now argue the opposite case. Assume your answer above is wrong. What's the strongest case against your recommendation? What evidence would change your mind? What assumption is doing the most work in your reasoning, and what happens if it doesn't hold?
```

> *Why it matters: models are trained to be agreeable. They'll confirm your priors unless you actively force them not to. This single follow-up is the closest thing to a free second opinion you'll find anywhere.*

SCENARIO · GP ABOUT TO RECOMMEND A DEAL

"I'm about to recommend this investment to my IC. Steelman the case against it. What would a sceptical partner — the one who's been right on the last three contested deals — say is wrong with this thesis? Be specific. Cite which assumption fails first."

Iterate, don't restart.

The biggest mistake non-technical users make: they get a mediocre answer and start a new chat. The right move is to refine in the same conversation. The model already has the context — give it the correction and ask again.

LAZY VERSION

```
[Closes the chat. Opens a new one.
Rewrites the prompt from scratch.]
```

SHARP VERSION

```
Good, but two corrections. (1) You over-
weighted the revenue concentration risk —
the contracts are 3-year so the
concentration is less acute than it looks
at first glance. (2) You missed that the
founder previously sold a similar business
in 2018; that's relevant context. Redo the
assessment with those two corrections.
```

> *Why it matters: every iteration in a conversation makes the answer sharper, because the model carries the context forward. Starting over throws away the work you've both already done.*

SCENARIO · ANYONE, ON ANY COMPLEX PROMPT

Treat the chat as a working session, not a vending machine. Push back, correct, add context, ask for the same answer in a different format, ask it to compress, ask it to expand. The fifth iteration of a prompt is almost always materially better than the first.

— THE TAKEAWAY

A 30-second checklist *for any high-stakes prompt.*

Before you hit send on anything that matters — a memo, a briefing, a decision-support analysis — run through this list. If you've covered seven or more, the answer you get back will be in a different league than what most people get from the same model.

THE PRE-FLIGHT CHECK

01 I've told the AI who it is and who I am.

02 I've told it what to do when it isn't sure.

03 I've pasted in the actual source material — not a paraphrase.

04 I've sequenced the steps it should take, in order.

05 I've shown it an example of what good looks like.

06 I've restated the ask in one line right before the data.

07 I've told it to reason out loud before giving the answer.

08 I've specified the format I want the answer in.

09 I'll ask it to argue against its own answer before I act on it.

10 I'll refine in the same chat, not start over.

— WHERE THIS CAME FROM

Information architecture is the *real bottleneck* in modern leadership.

The reason a sharp prompt produces a sharp answer is the same reason a well-run board meeting produces a sharp decision: structured inputs, clear sequencing, explicit handling of uncertainty, and the discipline to pressure-test the answer before acting on it. Most leadership failures aren't intelligence failures — they're information architecture failures.

That's the thesis behind **BoardLens**, the post-investment intelligence platform we're building for funds and boards. If you're a GP, board director, or operator wrestling with how your portfolio companies actually surface the signals that matter, I'd love to talk. We're taking a small number of design partner conversations this quarter — reply to any issue of *The AI Leadership Edge* and I'll be in touch.

Raffaela Rein

Founder, BoardLens · Writer, The AI Leadership Edge

Found this useful? Forward it. Every new subscriber via referral gets the next playbook before it goes public.