

Prompting for Leaders

Cheat Sheet

Quick-reference guide for getting better outputs from any AI tool

THE FIVE PROMPT ELEMENTS

1 Task

What do you want the AI to do? Be specific — write, summarise, compare, analyse, draft, extract, rewrite.

2 Context

What background does the AI need? Industry, company, project, recent events, specific constraints.

3 Audience

Who is this for? Board, executives, team, clients, investors — each needs different tone and depth.

4 Constraints

What limits or rules apply? Length, format, what to include or avoid, tone, specific requirements.

5 Format

How should the output look? Bullet points, numbered list, table, short paragraph, memo, slide outline.

WEAK VS STRONG PROMPTS

X WEAK PROMPT

"Summarize this."

Missing task specifics, audience, format, and length. The AI has nothing to calibrate against.

✓ STRONG PROMPT

"Summarize this report in 5 bullet points for a board audience. Highlight the 2 most material risks and suggest one recommended action. Keep it under 150 words."

Has task + audience + constraints + format. The AI knows exactly what to produce.

X WEAK PROMPT

"Write an email about the project."

No audience, no context, no purpose, no constraints.

EXECUTIVE RULE

Confidence in tone is not evidence of reliability in content.

Treat every AI output as a draft requiring judgment — not a certified answer ready to act on.

THREE FAILURE MODES TO KNOW

Hallucination

AI invents or distorts facts while sounding confident and authoritative.

Bias

Output reflects skewed assumptions from training data — especially in people decisions.

Overconfidence

Uncertain information presented with professional polish — fluency masks the gaps.

BEFORE ACTING ON AI OUTPUT — ASK

Is it useful?

Does this actually answer what I needed — or did it answer something adjacent?

Is it accurate?

Are the facts, numbers, and claims something I can verify or stand behind?

Is it complete?

What is missing, understated, or assumed that the audience needs to know?

Is it safe to use?

Does sharing this create legal, privacy, or reputational risk?

HOW TO REFINE AI OUTPUT

- Restructure**
Turn into a one-page memo — separate decisions, assumptions, and recommended next steps.
- Prioritise**
Ask for the top 3 issues, most material risks, or what matters in the next 30 days only.
- Surface uncertainty**
Ask: "What are you unsure about in this response? What would you need to verify?"
- Reframe for audience**
Rewrite for board vs team vs client — each needs a different level of detail and framing.