

# A CEO's Guide to Using AI: Tools to Use, Traps to Avoid

*With practical prompts you can use immediately*

*by Rens van den Broek and David Longworth*



*This guide serves as an actionable framework for using AI to sharpen executive judgment while avoiding the traps that erode it.*

Introduction to Executive Decisionmaking + AI	<u>2</u>
The Two Stages of Executive AI Adoption	<u>3</u>
Five Traps and Five Tools	<u>4</u>
<i>The Five Traps Where AI Undermines Executive Leadership</i>	<u>4</u>
<i>The Five Tools Where AI Strengthens CEO Judgement</i>	<u>7</u>
Conclusion: What This Means	<u>10</u>

What makes executive decision-making hard is not only making choices, but knowing which choices matter most, which assumptions really drive outcomes, and what is signal versus organizational noise. Most leaders who become CEOs have developed sophisticated instincts for these distinctions over time, but even the best instincts have limits. It's hard to surface your own blind spots or challenge your own logic.

AI changes this in ways that are both powerful and dangerous.

Much has been written about what AI technologies to invest in and how to build the right operating model to scale them. Those questions matter. But a quieter, equally important question gets less attention:

How should CEOs and senior executives use AI in their own work, [especially on the high-stakes decisions, judgment calls, and moral trade-offs](#) that define the job?

Research on executive decision-making has long shown that senior leaders are primarily compensated for judgment under uncertainty. Yet as authority increases, time becomes scarcer, judgment harder, information more filtered, challenge rarer, and the consequences of error greater. Against that backdrop, generative AI has the potential to multiply leadership. Used well, it surfaces gaps in thinking, generates structured dissent, and can focus attention where it has the highest potential return. That can mean faster, better decisions and more time spent on what matters most.

However, the same tools can quietly erode the very judgment they are meant to support. AI does not simply make leaders faster or more efficient. Over time, it changes how they think, where they direct their attention, and when they take shortcuts. It offers something tempting to leaders who are short on time and long on judgment calls: fast answers to difficult problems, surface-level coherence and conviction, and psychological distance from the cost of being the sole decision-maker.

Most executives already use AI in some form. In ghSMART's recent research, 47% of executives ranked leadership effectiveness as more important a driver of AI returns than technical talent, workflow integration, or organizational culture combined. What we explore in this article is a double-click on that finding: What does it look like when a CEO uses AI well versus poorly? Where and how does AI strengthen executive judgment, and where does it do the opposite?

CEOs are not all starting from the same place with AI. In our work, most fall into one of two stages: first, getting comfortable using AI personally; second, learning how to apply judgment to when and how they use it. This distinction matters because the risks and opportunities differ in each stage.

### **Stage 1: Get Comfortable Using AI Every Day**

The first step is straightforward: make AI part of your daily workflow. While most executives are now using AI, their engagement remains shallow, just 1.5 hours per week on average according to the [National Bureau of Economic Research](#). If your use looks similar, the real risk is falling behind while others build fluency.

If you have not yet done so, get access to a leading enterprise-grade AI assistant (such as Claude, ChatGPT, or Gemini, depending on what your IT team has approved) and start using it every day. Your early prompts do not need to be sophisticated: “Summarize this board deck.” “Draft talking points for my investor call.” “What questions should I expect on this earnings call?” The goal is not mastery; it is to build familiarity with how AI responds, where it surprises you, and where it falls short.

You need enough repetitions to develop intuition for what AI is actually good at versus what it merely sounds good at.

### **Stage 2: Apply Judgment to When and How You Use AI to Enhance Judgment**

Once you have that baseline fluency, the questions become harder and more consequential.

This is the stage where role-modeling matters. AI adoption operates with a visibility multiplier: when a CEO uses AI openly, the impact is not additive but multiplicative. When a leader asks AI to help them prepare for meetings, summarize how they have spent their time, or draft a first pass at communications, they do more than gain efficiency. They normalize experimentation and signal that fumbling with a new tool is not a liability and make it psychologically safe for others to follow.

A recent [BCG report](#) found that C-level executives who are deeply engaged with AI make their organizations 12 times more likely to be among the top 5% of AI performers. For CEOs, this means openly experimenting with AI and setting the stage publicly on where AI should and should not be used.

By contrast, CEOs who experiment privately while publicly demanding faster adoption from the rest of the team create a familiar pattern: middle managers express enthusiasm upward and skepticism downward, frontline employees comply without conviction, and the organization generates pilot theatre, an activity that resembles transformation but produces none of its results.

For CEOs who are past the “getting started” phase, the question is no longer whether to use AI, but how. In our work with early-adopter CEOs, we have found that AI functions simultaneously as a set of powerful cognitive tools and seductive traps.

The tools sharpen your judgment. The traps erode it — and the tricky part is that the **traps often feel like tools until the damage is done.**

What follows are five traps to avoid and five tools to embrace. We lead with the traps because they are where the real danger lies, and because the damage of falling into them is far greater than the upside of using AI well elsewhere.

One overarching note: everything that follows assumes AI has access to high-quality, relevant data. A frontier model working from public information alone is useful but limited. The real power comes when AI is connected to enterprise data, meeting notes, internal communications, and strategy documents. Without that, you are getting generic pattern-matching instead of contextual insight. If you have not yet worked with your team to figure out what data AI can safely access, that is your first priority after getting the subscription.

### The Five Traps (Where AI Undermines Executive Leadership)

**The most dangerous uses of AI are not obviously reckless.** They are tempting precisely because they offer relief from discomfort, create psychological distance from difficult trade-offs, and provide speed under time pressure.

The central risk is erosion: erosion of authority and trust when decisions feel outsourced, and erosion of judgment when AI takes over cognitive tasks it shouldn't. This risk doesn't appear all at once, it builds over time as leaders pass more decision-making to AI and the wrong habits take hold.



#### **1. The Silver Tongue - *Treating AI-generated coherence as truth.***

AI excels at producing coherent narratives from incomplete information. When facing ambiguous market signals or organizational tensions, a CEO might see AI synthesis as a way to resolve uncertainty and move forward with confidence. The output reads well. The logic flows. It feels like an answer.

That is exactly the danger. AI is exceptionally skilled at creating coherence, but coherence is not the same as truth. The risk is premature closure: weak but meaningful signals are overridden by patterns

that are legible and well-articulated. CEOs stop sitting with productive ambiguity, which is often where genuine insight emerges.

Many of the most consequential CEO decisions are meaning-making challenges, not analytical problems: What does this trend mean for us? How should we interpret this failure? What does our culture actually value? AI cannot answer these questions. It can only make them appear answered, which is far more dangerous than leaving them open.



## **2. The Autopilot - *Using AI as a substitute for judgment.***

AI offers speed and confidence. When facing a difficult trade-off, it is tempting to ask AI to analyze options, weigh variables, and surface a recommendation. The CEO reviews it, agrees it makes sense, and moves forward. It feels like augmentation, and it feels efficient.

Over time, it is something else: judgment atrophy. Each time you outsource the reasoning process — the actual weighing of competing priorities, not just the analysis — the muscle that leadership requires gets slightly weaker. The gap between what the role demands and what you can deliver on your own widens.

Using AI to challenge a decision the CEO has already formed is healthy. Using it to form the decision is where the atrophy takes hold. By the time the organization faces a crisis that demands decisive human judgment under conditions AI cannot model, the muscle may already be thinner than the moment requires.

The CEO (and by extension, organization) has, in effect, been training itself to defer rather than decide.



## **3. The Empty Robe - *Using AI for decisions that require moral judgment.***

Decisions with significant human impact are emotionally and reputationally costly. When facing layoffs, plant closures, or benefit reductions, it can feel reassuring to run scenarios through AI and let it structure the options. The distance can feel like objectivity.

AI has no inherent values, no lived experience of consequences, and no capacity for moral responsibility. It can optimize for defined parameters, but it cannot determine which parameters should matter or how competing values should be weighed.

When you use AI to inform decisions that hurt real people, the distance it creates is not objectivity. It is closer to an abdication of moral responsibility. Think of a judge's robe draped over an empty chair: all the trappings of authority, none of the human presence that gives authority meaning. Moral authority cannot be delegated to an algorithm; leadership requires owning the full weight of morally complex choices, publicly and completely.



#### **4. The Comfort Blanket - *Using AI to avoid discomfort rather than engage with it.***

AI can draft performance feedback, construct difficult messages, and articulate sensitive decisions with precision and tact. That can be a real help. But it can also become a way to avoid the emotional work of leadership.

The test is simple: Am I using AI to communicate better, or to feel more comfortable? If the answer is comfort, that is a signal to put the tool down and do the hard thing yourself.

People can detect, with remarkable accuracy, when a leader has not done the emotional labor themselves. The words might be correct, the tone might be professional, but something essential is missing: presence. Leadership is not transmitted through polished language. It is transmitted through the willingness to sit with discomfort, to own the difficulty of what is being said, and to remain present when things are hard. When leaders delegate this work, they do not save time. They train the organization to do the same: to avoid rather than engage.



#### **5. The Alibi - *Using AI as cover for accountability.***

AI feels objective. In a contentious restructuring or resource allocation, it can be tempting to reference AI analysis as a way to reduce personal exposure. “The analysis suggested” or “the model recommended” feels safer than “I decided.”

The moment people believe “the model decided,” authority erodes quickly. The team stops debating the substance of the decision and starts questioning the inputs. Then they stop trusting the CEO’s judgment and begin looking for ways to influence the system. Eventually, ownership diffuses entirely. Organizations follow accountability, not optimization. When AI becomes the perceived source of decisions, responsibility becomes impossible to locate.

There is also a practical problem. AI is not a spreadsheet. With a spreadsheet, you can go back and audit every assumption, trace every formula, and reconstruct exactly how a recommendation was derived. With AI, every prompt can produce a different response. Without deliberately capturing your inputs and AI’s outputs, there is no forensic trail. That matters for governance, for your board, and for your own learning. When AI informs a consequential decision, prompts and responses should be documented so the reasoning can be reconstructed later.

The problem is not that the CEO used AI. It is that the CEO positioned the decision in a way that made leadership appear delegated and left no trail to learn from.

## The Five Tools (Where AI Strengthens CEO Judgment)

If the traps are about AI replacing what should remain human, the tools are about AI amplifying what is already there.

The best uses of AI do not make the CEO less essential. They make the CEO more effective by surfacing what would otherwise remain hidden, generating challenges that would otherwise carry political cost, and making implicit reasoning explicit so it can be examined, improved, and shared.

The pattern across all five tools is the same: the CEO does the thinking. AI does the pressure-testing. The judgment stays human, and the rigor gets a boost.



### 1. The Sparring Partner - *Using AI to test your own logic.*


Every significant business decision rests on a small number of critical assumptions. Get those right, and details tend to work themselves out. Get them wrong, and execution excellence will not save you. The risk is that assumptions become invisible precisely because you treat them as facts.

AI is most valuable here after leadership has formed a view. The goal is to subject your own judgment to disciplined scrutiny.

#### Example prompts:

- *We are planning to consolidate three business units into one over the next 18 months. List the assumptions embedded in this plan that we are treating as facts rather than projections.*
- *Our Asia expansion is scheduled for Q3. What would have to be true for this strategy to fail within 18 months? Give me the failure modes we have not yet examined.*

This mirrors decades of research on pre-mortems and debiasing techniques, but with far less organizational friction. The value is in forcing the CEO to confront blind spots before the organization absorbs and acts on them.



### 2. The Red Team - *Using AI as a source of structured dissent.*

The Sparring Partner tests your logic. The Red Team gives voice to what people will not say out loud.


In most organizations, a CEO's draft decision receives careful review. Teams check the numbers, legal vets the language, and finance confirms the projections. What rarely happens is someone saying,

“Here is why this entire plan might fail”. Doing so implies doubt, challenges momentum, and can feel career-limiting.

AI has no career and no fear of appearing disloyal. That is what makes it useful here.

#### **Example prompts:**

- *Take the role of an activist investor who will acquire 8% of our company in 18 months. Review our current strategy and capital allocation. What concerns would you raise in a public letter to shareholders?*
- *Argue the strongest possible case against this decision. Do not hedge or provide balance. Be adversarial. The advantage is access to perspectives that organizational gravity prevents from surfacing. There is no status cost to dissent when the dissent comes from a machine.*



### **3. The Seismograph - Using AI as a pattern detector across fragmented signals.**

CEOs accumulate information the way rivers collect tributaries: slowly, from many directions, never all at once. A concern raised in Monday’s leadership call. A customer comment buried in a quarterly review. Three executives independently using the word “friction” to describe different problems. None of these signals arrives with a flag saying it is connected to the others.

Human memory excels at storing discrete events and poor at synthesizing patterns that emerge over weeks or months. By the time a pattern becomes obvious, the organization has often already committed to a path that is harder to reverse.

AI can help here, but only if you feed it real fragments: meeting notes, leadership conversations, email summaries, internal survey comments. With enough input over time, it can surface weak patterns before they harden into reality.

#### **Example prompts:**

- *After summarizing a month of leadership team conversations, ask: What themes appear repeatedly but never reach resolution? What topics generate discussion without producing decisions?*
- *Compare internal communications from Q1 and Q4 and ask: Which words have become more common? Which have disappeared? What might that reveal about culture or priorities shifting beneath the surface?*

Research on organizational culture consistently shows that drift appears in language before it appears in results. The Seismograph helps you see tremors before the earthquake.



#### **4. The X-Ray - Using AI to see past the polished surface.**

Every experienced CEO knows the danger of taking things at face value. The candidate who interviews brilliantly but leads poorly. The strategy presentation that is flawless on paper but hides a fatal assumption. The internal narrative that everyone repeats, but nobody has actually tested.

This is particularly acute with internal leaders whom the CEO knows well. Familiarity can masquerade as insight. However, what they truly know is a version shaped by proximity, goodwill, and organizational narrative. The same dynamic applies to strategy reviews: are we seeing what is really there, or what has been prepared for us to see?

AI will not make judgment calls for you, but it can help you identify where to probe.

##### **Example prompts:**

- *I have interview notes from three CFO finalists. Compare their narratives. Is any relying too heavily on rehearsed success stories? Are they avoiding discussion of specific types of failure?*
- *Here is the strategy deck my team presented. What assumptions are embedded in this that are never explicitly stated? Where does the logic depend on things going right that we have not tested?*

##### **You can even turn the X-Ray on yourself:**

- *Here is the all-hands email I am about to send about our restructuring. Read it as a skeptical mid-level manager who has seen two reorganizations in three years. What am I intending to communicate versus what will be perceived?*

The value here is attention management. AI helps the CEO spend scarce time probing the right risks rather than confirming polished surfaces.



#### **5. The Decoder Ring - Using AI to externalize decision logic.**

The most dangerous decisions CEOs make are not the ones that fail. They are the ones that succeed for the wrong reasons, where good outcomes mask flawed judgment and the organization never learns what really matters.

AI creates a way to make visible not just what was decided but how it was decided. The goal is to make judgment legible, transmissible, and improvable over time.

##### **Example prompts:**

- *After major decisions, prompt: I just chose option X over option Y. What principles or thresholds can you infer from this choice about how I weigh competing priorities?*

- Periodically ask: *Based on my last 20 decisions, when do I prioritize speed over consensus? When do I override financial metrics in favor of strategic intuition? What patterns emerge?*

You can then ask AI to draft a decision framework your leadership team can reference when you are not in the room. It will not replace you, but it can create a useful proxy for how you think, allowing the organization to move faster without constant escalation.


## WHAT THIS MEANS

If a decision requires courage, legitimacy, or moral ownership, AI can inform it. AI cannot decide it. AI can support cognition. It cannot replace accountability, authority, or human presence.

The unifying questions across all ten items above are simple: Will this use of AI make you and your organization smarter over time, or dumber? Are you using it to hone your edge, or to replace it? The CEOs who will use AI most effectively are those who see more clearly without mistaking coherence for truth, who challenge themselves more rigorously without abdicating responsibility for outcomes, and who scale their judgment without diluting the authority that makes leadership possible.



**Rens van den Broek** is a partner at ghSMART and Chief Growth Officer; he leads the firm's research on AI and leadership.



**David Longworth** is a principal at ghSMART and is leading firmwide adoption of AI.