

Using AI in the Decision Cycle

A practical guide to AI as a thinking partner

Using AI in the Decision Cycle © 2026 John Bornholt. Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0), creativecommons.org/licenses/by/4.0. You may copy, adapt and share this tool, including commercially, if you credit: "Adapted from Using AI in the Decision Cycle by John Bornholt, from LIFE: Our Uncertain Experiment, uncertaindecisions.com". Quotations and cited works belong to their authors and are not covered by this licence.

AI is a powerful thinking partner and a poor decision-maker. It can expand your options, challenge your assumptions, and identify considerations you might miss. It cannot weigh what matters to you, understand your full context, or carry accountability for the outcome. The decision and its consequences are all yours. Used well, AI helps to make a thorough decision process faster and more robust. Used poorly, it adds noise, creates false confidence, and outsources thinking that should stay with you.

PART 1 What AI is - and what it is not

The AI tools available today (large language models, for example ChatGPT, Claude and Gemini) are pattern-recognition systems trained on very large amounts of text. They are remarkably good at generating coherent, plausible prose, at summarising and reorganising information, at generating alternatives, and at explaining positions from multiple angles. They have no understanding of your specific situation beyond what you tell them. They have no stake in the outcome, and no consistent view. The same question asked differently can produce different answers.

This last point is worth understanding. Research from UCL and Oxford (Doshi et al., 2025) found that individual AI evaluations of strategic decisions are often inconsistent and biased. However, after aggregating evaluations across multiple prompts, roles, and phrasings, AI often produces results that resemble those of human experts. The practical implication here is to treat any single AI response as a starting point, not a conclusion. Ask the same question from different angles. Where the answers converge, there is probably something useful. Where they diverge, that divergence is itself informative.

PART 2 Where AI can help across the Decision Cycle

The table below describes the most useful role for AI at each step. The column describes what to ask for, not what to accept without scrutiny.

Step	Name	Where AI can help
1	Frame the Decision	Generate alternative framings of the presenting question. Challenge assumptions buried in the current frame. Ask: “What question might I be avoiding?”
2	Explore the Evidence	Identify possible knowledge gaps. Summarise background on unfamiliar topics (verify before relying on). Describe what typically happens in comparable situations (the outside view).
3	Challenge Your Thinking	Make the strongest case for supporting the option you are dismissing. List the cognitive biases that might be active. Generate disconfirming evidence you may not have looked for.
4	Weigh the Options	Brainstorm additional options. Identify options not yet on the table. Stress-test each option: “What is the strongest case against this?”
5	Assess the Risks	Identify risks not yet considered. Describe second-order outcomes from each option. Generate failure scenarios for the preferred option.
6	Test Your Values	Express the values trade-off in plain language. Describe how different value frameworks would evaluate this decision.
7	Decide and Commit	Summarise the case for and against each option. Identify what remains unresolved. Play the devil’s advocate against your intended choice. Draft the decision statement. Identify who needs to know and suggest what they need to hear.
8	Do and Deliver	Help build a first-action plan. Identify possible obstacles and responses in advance.
9	Reflect and Learn	Structure the after-action review. Identify patterns across multiple decisions when viewed as a whole.

PART 3 How to use AI well - five prompting principles

1 Give it full context.

AI cannot help with what it does not know. Describe the decision, the constraints, your current thinking, and what you are trying to work out. Vague input produces vague output.

Example: “I am deciding whether to [specific decision]. The constraints are [X, Y]. I am currently favouring [option] because [reason]. What am I missing?”

2 Ask it to disagree with you.

AI is usually agreeable. Instructing it to challenge, critique, or find problems will produce more useful output than asking for a balanced view.

Example: “What is the strongest case against the option I am preferring? What would a well-informed critic say about my reasoning?”

3 Ask from multiple angles.

Do not rely on a single response. Rephrase the question, assign it a different role, or approach it from the opposite direction. Where responses converge, the case is stronger.

Example: “Now argue the opposite position.” / “How would someone with [different background] see this?” / “What would the strongest counterargument be?”

4 Ask what you might be missing.

AI has been trained on a vast range of material and can often identify considerations, risks, or options outside your current knowledge or frame of reference.

Example: “What considerations relevant to this kind of decision am I not yet addressing? What do people in this situation typically overlook?”

5 Generate your own list before asking AI.

If you ask AI first, its output will anchor your thinking. Generate your own options, risks, or framings without AI assistance. Then use AI to check, extend, or challenge what you have done.

Example: Draft your own analysis first. Then: “Here is my current thinking [paste]. What have I missed? What would you add or challenge?”

PART 4 What to be cautious about

Hallucination. AI can state incorrect facts with complete confidence. It does not know what it does not know. Any factual claim made by AI (e.g. statistics, dates, research findings, legal or regulatory details) must be verified independently before being relied on. Do not treat AI output as a source of facts. Treat it as a source of directions to investigate.

Sycophancy. AI is trained to be helpful and agreeable, which means it often validates your framing and confirms your instincts unless told otherwise. The more you share your existing view before asking for input, the more likely the response is to reinforce it. Use Principle 2 above, and ask for disagreement to counteract this.

Confidentiality. Information shared with commercial AI tools may be used for training or stored in ways outside your control. Do not share sensitive personal information, commercially confidential details, information about identifiable third parties, or anything that would be damaging if disclosed. Where confidentiality matters, describe the situation in general terms without identifying details.

Anchoring. The first option, frame, or risk that AI suggests can anchor your thinking, even when you consciously try to consider others. Generate your own list independently before asking AI, then compare. If AI’s first response dominates your subsequent thinking, that is a sign you asked too early in the process.

Over-reliance. If AI does your thinking, then it is not your thinking. The Decision Cycle is designed to produce decisions that you own, that you understand, and that you can account for. Using AI to replace steps, rather than to enrich them, undermines that. Here’s a useful test to apply. Could you explain your reasoning to someone who asked, without referring to what AI said? If not, then you may have delegated too much thinking.

People make two opposite mistakes with advice from machines: accepting it without checking, and dismissing a system that is working well (Parasuraman and Riley, 1997). Some people abandon an algorithm after a single error (Dietvorst et al., 2015), while others trust advice more because it came from an algorithm (Logg et al., 2019). In other words, we are poor judges of how much we are trusting a machine. One question helps: does this decision rely on the AI knowing and understanding my particular circumstances? You can describe your circumstances to a machine, but that does not mean it understands them.

PART 5 Who is accountable?

AI has no stake in your decision. It will not live with the consequences. It will not answer to those affected. It will not be responsible for what follows. You will. And this matters, because the more consequential the decision, the more important it is that the reasoning is yours. It can be informed and enriched by AI, but not produced by it. Use AI as you would use a knowledgeable colleague who is not accountable for the outcome.

See also: Gathering Information Well (Step 2), Bias Awareness Checklist (Step 3), Generating Options (Step 4).

Sources and references

GENERATIVE AI AND DECISION-MAKING

Doshi, A. R., Bell, J. J., Mirzayev, E., & Vanneste, B. S. (2025). Generative artificial intelligence and evaluating strategic decisions. *Strategic Management Journal*, 46(3), 583-610.

Directly cited in this guide. Found that individual AI evaluations of strategic decisions are often inconsistent and biased, but that aggregating evaluations across multiple LLMs, prompts, and roles can produce rankings resembling those of human experts. Supports the guidance to probe from multiple angles rather than rely on a single AI response.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM FAccT Conference*, 610-623.

A foundational paper on the limitations of large language models, including the risk that fluent, confident output may not reflect accurate or grounded understanding. Relevant to the caution on hallucination in this guide.

COGNITIVE BIASES RELEVANT TO AI USE

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

Relevant to the anchoring and sycophancy cautions - the tendency to over-weight the first piece of information encountered, and the difficulty of challenging an initial framing once it has been accepted.

TRUSTING A MACHINE

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230-253.

Set out the two opposite failures in how people handle automated systems: over-reliance and neglect.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114-126.

Found that people abandon an algorithm faster than a person who makes the same mistake.

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90-103.

Found the opposite tendency: people often follow advice more readily when they believe it came from an algorithm.