

Before You Commit

How to think clearly through one decision that actually matters

Most bad executive decisions don't fail from missing data. They fail because the question was framed too narrowly, the information was corrupted by the people delivering it, the leader couldn't separate the decision from their identity, or a reversible bet was treated as irreversible — or worse, the other way around.

This guide is structured for one purpose: **improving the quality of a single high-stakes decision**. It is not a worksheet. It is not a checklist. It is a structured walk-through of the nine moves the best decision-makers make before they commit — and the traps the rest of us fall into when we don't.

Use the dashboard on the next page to navigate. For a fast decision, the dashboard alone may be enough. For a decision you'll live with for a decade, do the work.

Compiled by **Ratipcan Uysal**

ONE INSTRUCTION BEFORE YOU BEGIN

Write the decision you're facing in a single sentence. If you can't get it into one sentence, you don't yet know what you're deciding — and no framework will help. The clarity of that sentence is the single best leading indicator of the quality of the decision that follows.

One Page. Five Minutes. Real Decision.

If you only have a few minutes, do these five things. The rest of Before You Commit is depth behind these five questions.

1. THE SENTENCE

Write the decision in one sentence. Then write the question *behind* the question. Most CEOs are solving the second one without noticing.

2. THE DOOR

One-way door (irreversible) or two-way door (reversible)? Type 1 deserves weeks; Type 2 deserves hours. Match the process to the stakes — see Stage 02.

3. THE THIRD OPTION

If you have only two options, you haven't widened enough. Generate at least three genuinely distinct alternatives before evaluating any of them.

4. THE PRE-MORTEM

It is 18 months from now. This decision has failed. What does the autopsy say? Write it. Then prevent the top three causes before committing.

5. THE TRIPWIRE

What specific event would tell you this decision is going wrong? Set the metric and the review date *before* you commit — never after.

AND THE HONEST ONE

What are you protecting by making this decision the way you want to? Your ego, your story, someone's approval? Name it. Then decide anyway, eyes open.

BEFORE YOU PROCEED — TRIAGE BY ARCHETYPE

Different decisions fail in different ways. Identify your archetype first — then weight the stages and use the process that fits.

ARCHETYPE	SIGNAL	PROCESS	WATCH-OUT	STAGES
People	Hiring, firing, promotion, partnership	Full guide. Pre-mortem on the person, not just the role. Stage 08 first.	<i>Identity overrides evidence; acting too late</i>	5 • 6 • 8 • 9
Strategic Bet	Pivots, new markets, big product bets	WRAP. Red-team the thesis. Pre-mortem. Adversarial Thinking.	<i>Narrow framing; overconfidence; analysis paralysis</i>	1 • 3 • 4 • 6
Capital Allocation	Acquisitions, major spend, equity raises	Outside view first. Audit advisor incentives. Asymmetric Bet Sizing.	<i>Deal momentum; base-rate neglect; advisor incentives</i>	2 • 4 • 5 • 7
Crisis Response	Time-bound, high-pressure, public	OODA. Pre-assigned roles and runbooks. Use the time you have.	<i>Action bias; communication debt; emotion as strategy</i>	2 • 8 • 9
Reversible Experiment	Pricing, marketing, small product bets	Classify only. Ooch. Set tripwire. Skip the rest.	<i>Over-processing Type 2; pilot-to-trap escalation</i>	2 only
Innovation Bet	Long-horizon, high-uncertainty, no precedent	Real options: small staged bets. Keep doors open. Optionality card.	<i>Zombie projects; committing before signal; sunk cost lock-in</i>	2 • 3 • 6
Org Design	Restructuring, reporting lines, role elimination	Multi-Frame first. Map winners and losers. Chesterton's Fence before touching any structure.	<i>Announced early, decided late; Cobra Effect on removed constraints</i>	1 • 5 • 7 • 9
Exit / Kill	Product, market, or partnership discontinuation	Founder's question: would you start this today? Name the kill criterion cold.	<i>Sunk cost identity; killing too late; mistaking pivot for exit</i>	1 • 2 • 8
Value Trade-off	Competing priorities, significant human impact	Stakeholder map. Name the values explicitly. Stage 08 is the real work.	<i>Hidden value conflicts; false consensus; ethics as decoration</i>	7 • 8 • 9

CONTENTS

The nine moves

Each stage is one page. Before You Commit is designed to be read in 30 minutes — or referenced in 3, depending on what the moment requires.

-
- | | | |
|-----------|--------------|--|
| 01 | Frame | <i>What are you actually deciding?</i> |
|-----------|--------------|--|
-
- | | | |
|-----------|-----------------|--|
| 02 | Classify | <i>One-way doors and two-way doors</i> |
|-----------|-----------------|--|
-
- | | | |
|-----------|--------------|-------------------------------|
| 03 | Widen | <i>Escape the binary trap</i> |
|-----------|--------------|-------------------------------|
-
- | | | |
|-----------|-------------|--|
| 04 | Test | <i>Confronting the data you don't want</i> |
|-----------|-------------|--|
-
- | | | |
|-----------|---------------|--|
| 05 | Filter | <i>Information hygiene at altitude</i> |
|-----------|---------------|--|
-
- | | | |
|-----------|-------------------|--|
| 06 | Pre-Mortem | <i>Imagining the failure in detail</i> |
|-----------|-------------------|--|
-
- | | | |
|-----------|----------------|---|
| 07 | Surface | <i>Engineer dissent before you decide</i> |
|-----------|----------------|---|
-
- | | | |
|-----------|-----------------|---------------------------------------|
| 08 | Confront | <i>What you are really protecting</i> |
|-----------|-----------------|---------------------------------------|
-
- | | | |
|-----------|---------------|---------------------------------------|
| 09 | Commit | <i>Decide, communicate, calibrate</i> |
|-----------|---------------|---------------------------------------|
-

THE REFERENCE

- 100** **100 Questions to Ask Before Deciding Anything** — organized by stage

APPENDICES

- A** **The Biases That Sink CEO Decisions** — Twelve failure modes, with real stories
- B** **The Decision Journal** — Pre-commitment template
- C** **When NOT To Use This Guide** — Process isn't free
- D** **Self-Calibration** — Getting better, decision by decision
- E** **Question Techniques From the Literature** — 37 techniques — Hero–Dragon–Quest, Adversarial Thinking, Optionality, Asymmetric Bet Sizing, Inversion, Local vs. Global Maxima, Weighted Criteria & more

- F Key Insights from the Source Library** — Sharpest ideas from Decisive, How to Decide, Solvable, Berger & Farnam Street
- G The Reading List** — Curated — books, essays, websites
- H All Sources** — Every book, article, and essay drawn on in this guide

01

Frame

What are you actually deciding?

The single most common executive error is solving the wrong problem with great rigor. The framing you inherit — from a board member, an advisor, a Slack thread, your own anxiety — is almost never the framing that serves you. Before doing anything else, surface the framing and challenge it.

Good framing has a specific signature: it states the decision in a single sentence, at the right altitude, with the right scope, and without baking in the answer. Most framings fail one or more of those tests on first inspection.

"If I had an hour to solve a problem I'd spend 55 minutes thinking about the problem and 5 minutes thinking about solutions." — attributed to Einstein

The framing diagnostic

1. State the decision in one sentence. Now state it three more times, each broader than the last. Which framing is actually the right altitude? *Example: "Should we fire Marcus?" → "Do we have a performance problem or a role-fit problem?" → "Is this team structured to win?" Each reframe opens a different set of options.*
2. What is the **question behind the question**? Why is this decision in front of you *now*? *Example: A board pushing for an immediate CFO hire — the real question behind it is often "do they trust the current financial controls?" Answering the surface question leaves the actual one untouched.*
3. If you do nothing — make no decision — what actually happens in 30, 90, and 365 days? (Inaction is a decision.) *Example: "Should we respond to the competitor's price cut?" — if you do nothing for 90 days and lose 3% market share, that answer tells you how urgent this actually is.*
4. Are you choosing **whether** to do something, or **how** to do it? These are different decisions and need different processes. *Example: "How do we enter the German market?" assumes the entry is decided. The whether question — "should we enter Germany at all?" — was never asked.*
5. Whose decision is this, really? If it's not yours alone, who else holds a veto, and have they been engaged? *Example: A CEO spending three weeks modeling a partnership deal, only to discover the board chair considers it outside the delegated authority. The framing step is also a mandate check.*
6. How reversible is this — a week, a quarter, never? (This question gets its own stage. Hold the answer.)
7. If a smart outsider walked in cold, would they frame it the way you have? Why or why not? *Example: Every CEO of a struggling company frames the core problem differently than a new hire in week two. The outsider's framing is almost always closer to the truth — and almost always more uncomfortable.*
8. **Is a solution already baked into the problem statement?** When someone says "we need better onboarding to improve retention" — the solution (onboarding) is already embedded. The question of whether retention is actually an onboarding problem never gets asked. Strip the solution out and restate the problem as a pure gap: what outcome are we actually trying to change? *Another example: "We need a new CRM" is a solution. "Our sales team doesn't have the information they need, when they need it" is the problem. Those have very different answers.*

The framing trap — your lens is limiting your options

1. **What is your default frame?** People? Process? Technology? Culture? Market? Whatever it is, it determines which causes you can see and which solutions you can reach for — without you noticing. Name it explicitly before diagnosing anything. *Example: A founder who built the company through sheer talent density will diagnose almost every problem as a hiring problem. A process-oriented operator will see the same symptoms as a workflow problem. Both are partially right. Neither can see the other's answer.*
2. **Apply at least two other lenses** to the same situation. A problem you've been framing as a people issue: what does it look like through a systems lens? A process lens? A market lens? The solutions you couldn't see before usually live in the frames you haven't tried.
3. **Is this the symptom or the problem?** The pain someone reports — wasted time, low morale, missed deadlines — is almost never the actual problem. Ask why the symptom exists, and keep going until you hit something structural. The strategic error is treating the effect as the cause. *Example: Engineers are leaving. The framing is "retention problem." Ask why three times: they're leaving → for better-paying roles → because the company hasn't raised in 18 months → because the Series B narrative isn't landing. The real problem is pitch strategy, not HR policy.*
4. **Where is this problem *not* occurring?** If sales are down in every region except one, the question isn't "why are sales down" — it's "what is that region doing differently?" Bright spots tell you something the problem itself cannot.
5. **Would solving this problem destroy something useful?** (Chesterton's Fence.) Before removing any constraint, rule, or process — understand why it exists. If you can't explain the function it serves, you're not ready to remove it. *Example: A new COO eliminates the weekly all-hands as "theater." Six months later, the information flow that held cross-functional alignment has silently collapsed. The meeting was annoying. It was also the only place three departments compared notes.*
6. **What are the second and third-order effects of this decision?** What are the effects of the effects? Solutions that fix the first-order problem while creating downstream consequences are the most expensive kind of wrong.
7. **Is this problem worth solving right now?** Organizations have a limited rate at which they can absorb change. Even a real problem, correctly diagnosed, can cause damage if solved at the wrong moment — when capacity is stretched, trust is low, or the team is already absorbing ten other changes. *Example: A restructuring that would have worked in January becomes destabilizing in October, four months after a difficult layoff. The problem was real. The timing was wrong. The distinction matters.*

THE ELEVATOR WITH MIRRORS

Tenants in a building complain the elevator is too slow. The obvious solution space: upgrade the motor, improve the algorithm, install a new elevator. All expensive, all addressing the stated problem. A building manager suggests something different: put mirrors next to the elevator. Complaints stop. Not because the elevator got faster — it didn't — but because people lose track of time when given something fascinating to look at. The frame was wrong. The problem wasn't speed; it was the experience of waiting. The frame you inherit determines the solution you can see. This story, from Thomas Wedell-Wedellsborg's *What's Your Problem?*, is the cleanest illustration of why framing matters more than analysis.

02

Classify

One-way doors and two-way doors

Not all decisions deserve the same process. Bezos's distinction is the cleanest frame yet articulated: **Type 1** decisions are one-way doors — irreversible or near-irreversible, with consequences that can't be undone. **Type 2** decisions are two-way doors — you can walk back through if you don't like what's on the other side.

The error most CEOs make is symmetric and expensive: they apply Type 1 caution to Type 2 decisions, and slow the company to a crawl. Or they apply Type 2 speed to Type 1 decisions, and blow the company up. Both errors look like good judgment from the inside.

"Most decisions should probably be made with somewhere around 70% of the information you wish you had. If you wait for 90%, you're being slow." — Bezos

Identify the door

1. If this decision turns out wrong, how long does it take to reverse — and at what cost in money, reputation, and relationships?
2. What's the cost of being **wrong** vs the cost of being **slow**? Whichever is bigger should dominate your process.
3. Am I treating a reversible decision as irreversible (over-engineering) or vice versa (under-investing)? Where exactly is the mismatch?
4. Is the decision **large** (consequential) or merely **loud** (emotionally charged but not very consequential)? These get confused constantly.

Then change your process accordingly

1. For Type 1: slow down. Get weeks, not hours. Get unfiltered input. Sleep on it. The decision is worth being late to be right.
2. For Type 2: speed up. Stop deliberating. Make it, ship it, learn from it. Most CEO over-thinking lives here — and it has a real cost.
3. If you're genuinely unsure: assume Type 2 unless you can specifically name what makes it Type 1. Most decisions feel more irreversible than they are.
4. **Organizational absorption rate.** Even a correct decision, well-timed in isolation, can cause damage if the organization is already absorbing too much change. What's the change-load right now? Is there capacity — cognitively, operationally, emotionally — to absorb this too?

THE TRAP — TREATING ALL DECISIONS AS LARGE

The most expensive CEO time is time spent on decisions that aren't yours, or that don't matter at the scale you're treating them at. If you're spending three days on a decision your VP could make in three hours, you have a classification problem, not a decision problem.

03

Widen

Escape the binary trap

If Stage 1 revealed that you're stuck in a yes/no, this is where you fix it. The cheapest way to improve a decision is to invent better options before evaluating any of them. Constraints sharpen creativity; widening at this stage is what prevents "the alternative we never considered" from becoming the post-mortem's first finding.

Standard to hit before evaluating: at least **three genuinely distinct options** — not three flavors of the same option.

Option-generation moves

1. **The vanishing-option test.** If your preferred option suddenly became illegal, what would you do instead? That answer is option B — usually better than you assumed.
2. **The AND, not OR move.** Can you do *both* of two competing options at smaller scale? Many "either/or" choices dissolve under this question.
3. **The opposite move (inversion).** What's the exact opposite of what you're planning? What would have to be true for that to be correct?
4. **The 10× and 1/10 move.** What would you do if you had to move 10× faster? With half the budget? Twice the budget? Each constraint generates a real option.
5. **The benchmark move.** Who has faced this exact decision before — in or out of your industry? What did the best ones do? What did the worst?
6. **The disowned option.** What's the option you've ruled out without examining? Often the most useful to surface, even if you ultimately reject it.

WHEN TO STOP WIDENING

Generating more options indefinitely is its own form of avoidance. Once you have three to five genuinely distinct alternatives — and you can articulate the logic behind each — stop widening and start testing. The art is in the transition, and the right moment is when option *quality* stops climbing.

04

Test

Confronting the data you don't want

By this stage, you have a framing, a classification, and several options. Now the question is: **what do you actually believe, and is it true?** Most executives stop testing once they've found enough evidence to support what they already wanted to do. The discipline here is to test the other way — actively, deliberately, against your own preference.

"What the human being is best at doing is interpreting all new information so that their prior conclusions remain intact." — Warren Buffett

Disconfirmation drills

1. **Name the load-bearing assumption.** What single belief, if false, would make this decision wrong? How confident are you, on what evidence? *Example: An acquisition thesis that rests on "their customer base will cross-sell to ours" — if retention rates for cross-sold customers are actually 40% lower, the entire model breaks. That assumption is load-bearing. Have you tested it — or just assumed it?*
2. **Run the disconfirming search.** Spend equal time looking for evidence *against* your preferred option as you spent finding evidence for it. Most CEOs spend 90/10. Fix the ratio. *Example: If you spent three hours building the case for entering a new market, spend three hours finding the best argument for why you shouldn't. Not a token scan — a genuine search. If you can't find any, you haven't looked.*
3. **Find the smart person who disagrees.** Not someone who'll soften it. Someone whose job, identity, or compensation doesn't depend on agreeing with you. *Example: Your CFO thinks this is a good deal. Your head of sales thinks it's a good deal. Now find the senior operator at a company that tried something similar and failed — and ask them why it didn't work. That conversation is worth more than the model.*
4. **Check the reference class.** What happened to other companies that made this decision in similar circumstances? Don't reason from your own narrative — look at the base rate. *Example: You're considering a major ERP implementation. The base rate on large ERP projects: over 50% come in more than 180% over budget and significantly late. Your situation feels different. It always does from the inside. Check the number before discarding it.*
5. **Ooch before committing.** Can you run a small experiment that produces real information before the big bet? Don't predict, prototype. *Example: Before signing a three-year exclusive distribution deal for a new market, run a six-week pilot through a non-exclusive partner. The cost is low; the information — real demand, real logistics, real friction — is irreplaceable. If the pilot is impossible, ask why, and whether that's a signal.*
6. **The five whys.** Why do you believe X? And why is *that*? Five levels down, do you still have a reason — or just a feeling? *Example: "I believe this hire is right." Why? "Strong interview." Why does that matter? "Past track record." Why trust that? "Reference checks were positive." Who gave the references? "He did." There it is — self-selected references are not a reference class. You're five levels down and the foundation just shifted.*

THE CONFIRMATION TELL

You'll know you're falling into confirmation bias when you notice yourself feeling *impatient* with disconfirming information rather than curious about it. That impatience is the tell. Slow down precisely there. The information you're rushing past is probably the information you need most.

05

Filter

Information hygiene at altitude

Here is the part most decision frameworks omit: **you are not deciding with neutral information.** Every fact reaching you has been filtered by someone whose career, comfort, or self-image depends on which conclusion you draw. The higher you climb, the more this is true.

At CEO altitude, your information environment is structurally corrupted — not by lies, but by selection, emphasis, framing, omission, and the silent editing people do before bringing things to you. A serious decision-maker treats this as the dominant problem, not a footnote.

"As you go up the org chart, the average level of incompetence at gathering facts goes up. People filter information for the boss." — Andy Grove

Audit the information getting to you

1. **Who has financial or career skin in this answer?** What direction would they want you to lean, and how might that be shaping what you've been told?
2. **What's the chain of edits?** Trace one critical fact back to its source. How many people did it pass through? Who could have softened, sharpened, or omitted what?
3. **What hasn't been brought to you?** List three things relevant to this decision that no one has mentioned. Ask why.
4. **Who would tell you something inconvenient?** Name them. If you can't — that is the problem. Fix the input system, not just this decision.
5. **Skip-level the data.** Talk directly to someone two or three levels below the people briefing you. The information is different, and usually better.
6. **Go to the source.** For a critical fact, talk to the customer, the engineer, the operator — not the person summarizing them. Even ten minutes of direct contact dissolves layers of distortion.

THE CEO INFORMATION TRAP

The most dangerous information is not what you don't know — it's what you *think* you know, that has quietly been shaped by the incentives of the people delivering it. Build a small number of relationships with people who are paid to disagree with you, or whose careers don't depend on you. They are the only ones who can tell you the truth at scale, and they are worth more than almost any other input you have access to.

06

Pre-Mortem

Imagining the failure in detail

Gary Klein invented the pre-mortem in the 1980s while studying firefighters and military commanders. Studies suggest it surfaces about 30% more reasons a project could fail than any other forecasting method. The mechanism: imagining a failure that has *already happened* unlocks specifics that imagining one that *might* happen does not. Klein's team also developed a companion technique — the **crystal ball** — for when you suspect you're already fixated on the wrong explanation: pretend a crystal ball has shown your current reading of the situation is wrong, without revealing why. Force yourself to generate an alternative. Do it twice. The explanations you reach are usually more accurate than the original.

Spend 20 minutes — alone or with your team — completing the pre-mortem in writing. Not in conversation. Writing is what forces specificity.

The pre-mortem protocol

1. **Set the scene.** It is 18 months from today. The decision you're about to make has failed catastrophically. The damage is real and visible.
2. **Write the autopsy.** What, specifically, went wrong? Don't generalize — name the people, moments, assumptions, missed signals.
3. **Identify the first domino.** What was the first thing that started going wrong? When could it have been caught?
4. **Name the warning signs.** What leading indicators of failure did you ignore, rationalize, or fail to have visibility into?
5. **Identify the external shocks.** What outside events made it worse? Were they truly unforeseeable — or just unconsidered?
6. **Locate your own contribution.** What did *you* do, or fail to do, that contributed? Be specific. This is where the intelligence lives.
7. **Reverse-engineer the fix.** For each cause you named, what would you have to put in place *now*, before deciding, to prevent it?
8. **The crystal ball.** Now pretend a crystal ball has shown you that your *current understanding of the situation* is also wrong — without revealing how. Generate two alternative explanations for what's happening. If they don't change your plan at all, you may be fixated.
9. **The four blind spots.** Run through the four categories of what you might have missed: (1) *Out of context* — what was hidden from your circumstances at the time? (2) *Out of mind* — what simply didn't come to mind? (3) *Lackluster* — what was crowded out because something else grabbed your attention first? (4) *Expired* — what option did you neglect because it didn't seem new or exciting? Each category surfaces a different kind of gap.

WHY THIS WORKS

Prospective hindsight — imagining a future event has already occurred — bypasses the optimism bias that distorts forward-looking forecasts. Your brain treats the failure as a fact to be explained, not a possibility to be argued against. The explanations are usually accurate. Take them seriously.

07

Surface

Engineer dissent before you decide

Most consequential decisions involve a board, an exec team, key investors, or a co-founder. The hard question is rarely "what do I think?" It's "have I designed the room so that what other people *actually* think can reach me?"

Genuine dissent does not arise spontaneously around powerful people. It has to be engineered. The CEOs who decide best build the social architecture of dissent *before* the decision is on the table — so when it matters, the room can tell them the truth.

"If there's anyone in this room who agrees with everything I just said, I want you to leave."
— Alfred P. Sloan, reportedly, before a major GM vote

Architecting dissent

1. **Solicit privately first.** Get individual written input from each key stakeholder *before* the group meets. The collective conversation overwrites what people would say alone. The Japanese call this *nemawashi* — literally, tending the roots before the tree is moved. The consensus you build in private is the only consensus that holds in public.
2. **Speak last.** If you state your preference first, the conversation will converge on it whether or not it's right. The most senior person should be the final voice, not the first.
3. **Assign the devil's advocate.** Designate someone — by name, in advance — to argue against the leading option. Make it a real role, not a performance. Rotate it.
4. **Ask: "what would make us not do this?"** Not "any concerns?" — that gets nothing. The specific question forces specific answers.
5. **Test for compliance vs commitment.** Disagree-and-commit only works if disagreement was actually voiced. Silent agreement isn't alignment — it's risk.
6. **Build a kitchen cabinet.** Maintain two or three people, inside or outside the company, who tell you what no one else will. The relationship has to predate the decision. Build it now if it doesn't exist. The three criteria: (1) they have no financial stake in your conclusions; (2) they have seen enough of your blind spots to name them; (3) they are willing to make you uncomfortable. One board member, one peer outside your industry, one person who has failed at something you're attempting.

THE DEEPER TRAP

If you've never gotten genuine pushback from your team on a major decision, the problem is not that your team agrees with you. The problem is that you've built a room where disagreement isn't safe. That's a deeper failure than any single decision — and fixing it produces compound returns. Start with one meeting where you explicitly invite, reward, and act on dissent. Word travels.

08

Confront

What you are really protecting

Most decision frameworks treat the CEO as a rational analyst processing inputs. Real CEOs decide while their identity, ego, reputation, sunk narrative, and self-image are all in play simultaneously. Pretending those forces aren't operating doesn't neutralize them — it just hides them.

The most consistent reason a CEO who knows the right answer doesn't take it is that the right answer threatens something about who they are or how they're seen. Naming that, honestly, is sometimes the entire decision.

"The first principle is that you must not fool yourself — and you are the easiest person to fool." — Richard Feynman

The honest interrogation

1. **What are you protecting?** By making this decision the way you want to, what stays intact? Your story about yourself? Your image with the board? A past commitment you'd have to walk back?
2. **Whose approval are you optimizing for?** Not abstractly — name them. Would you decide differently if their opinion didn't exist? If yes, that's information.
3. **What's the version of you that would make a different call?** Who would you have to become — even for ten minutes — to decide more clearly?
4. **What past commitment is this decision actually re-litigating?** If you publicly committed to X, is this decision really about X — or about not having to say you were wrong about X?
5. **How would you decide if you'd never said anything about this publicly?** The difference between that answer and your current one is the cost of your ego in this decision.
6. **Is fear shaping this?** Specifically: fear of looking weak, fear of disappointing someone, fear of being wrong, fear of irreversibility itself. Name the one operating most. Then ask whether it's a useful signal or a distortion.
7. **The Regret Minimization test.** Project yourself to age 80 and look back on this decision. Which choice would you regret more — acting, or not acting? Bezos used this explicitly when deciding to leave a secure job to start Amazon. It doesn't override analysis, but it cuts through the short-term noise that makes near things loom larger than they deserve.

THE QUIET ONE

The decision you keep *not* making is almost always a decision your identity is protecting you from. Firing the wrong executive. Killing the wrong product. Admitting the wrong strategy. The longer it sits, the more it costs — and the more it warps every adjacent decision. If something has been on your plate for months unmade, ask whether the problem is the analysis (rare) or the identity cost of the answer (usually).

09

Commit

Decide, communicate, calibrate

The decision is not the moment of choice. It is everything that happens after. Even excellent thinking falls apart at this stage: leaders relitigate the call in their head for weeks, fail to communicate it cleanly, miss the early signs it's going wrong, and don't capture what they actually believed when they decided — so they can't learn from it later.

Commit, communicate, calibrate. In that order. Before you walk out of the room.

Commit — close the loop

1. **Lock in the decision.** Write down what you've decided, why, and what you considered. Date it. The act of writing it down is what stops the relitigation.
2. **Set the tripwire.** What specific metric or event would tell you this is going wrong? Set the number, not the vibe, and put the review date on the calendar.
3. **Name your kill criterion.** The point at which you would reverse the decision — no matter how much has been invested. Decide it now, in cold blood.

Communicate — make it real

1. **Tell the disagreeers first.** The people who pushed back deserve to hear the decision directly, before it's announced. Otherwise you trade short-term comfort for long-term trust.
2. **Explain the reasoning, not just the conclusion.** People can disagree-and-commit to a reasoned call. They can't commit to an edict. The reasoning is the alignment.
3. **Say what would change your mind.** Make it explicit. This both invites useful new information and protects the decision from drive-by re-litigation.
4. **Stop selling the decision after it's made.** Over-justifying signals doubt. State it once, clearly, and move.

THE FINAL DISCIPLINE

When the review date arrives, separate **decision quality** from **outcome quality**. A good decision can produce a bad outcome — you got unlucky. A bad decision can produce a good outcome — you got lucky. Judging yourself on outcomes alone teaches the wrong lessons. Judge the decision on what you knew when you made it. Adjust the process, not just the conclusion.

THE REFERENCE

100 questions to ask before deciding anything

Distilled from Before You Commit and from how the best decision-makers actually think. Organized by stage. You will never use all of them on one decision — but for any decision worth making well, ten of these are the right ten.

01 • FRAME — what are you actually deciding?

Get the question right before answering anything.

1. Can you state the decision in a single sentence — without baking in the answer?
2. What is the question *behind* the question? Why is this on your plate now?
3. If you took no action, what specifically happens in 30, 90, and 365 days?
4. Are you choosing *whether* to do something, or *how* to do it?
5. Whose decision is this, really? Who else holds a veto?
6. What problem are you actually solving — and is it the right problem, or just the most painful one?
7. Am I hearing the symptom (the pain being felt) or the actual problem? Have I dug far enough to find the structural cause?
8. What is my default frame — people, process, technology, culture, market? What causes am I blind to because of it?
9. Would solving this problem destroy something useful I haven't noticed yet? (Can I explain why this constraint or rule exists before removing it?)
10. What are the second and third-order effects? The effects of the effects?
11. Is this problem worth solving *right now* — given what else the organization is absorbing?

02 • CLASSIFY — one-way doors and two-way doors

Match the process to the stakes.

12. Is this reversible — in a week, a quarter, never? At what cost?
13. Which is bigger here: the cost of being wrong, or the cost of being slow?
14. Am I over-engineering a Type 2 decision, or under-investing in a Type 1?
15. Is this decision *large* (consequential) or merely *loud* (emotional)?
16. What's the worst plausible outcome — and how recoverable is it?
17. If I'm wrong about this in 90 days, what would I do about it then?
18. Does the speed of this decision itself give me strategic advantage or disadvantage?
19. What's a defensible deadline for this — set in cold blood, not in the moment?
20. Is the decision actually mine to make alone — or am I borrowing authority that isn't mine?

03 • WIDEN — escape the binary trap

The cheapest way to improve a decision is to invent better options.

21. If my preferred option suddenly became illegal, what would I do instead?
22. Can I do *both* competing options at smaller scale?
23. What's the exact opposite of what I'm planning — and when would it be right?
24. What would I do if I had to move 10× faster? With half the budget? Twice it?
25. Who has faced this exact decision before? What did the best ones do?
26. What option have I ruled out without actually examining it?
27. What would the founder version of me do? What would the CEO version?

28. What's the option that scares me — and is the fear real or just unfamiliarity?
29. What would I do here if I knew I could not fail?

04 • TEST — confronting the data you don't want

Test against your preference, not in favor of it.

30. What single belief, if false, would make this decision wrong?
31. How confident am I in that belief — and on what specific, checkable evidence?
32. Have I spent equal time looking for disconfirming evidence as confirming evidence?
33. What's the smart person's strongest argument against my preferred option?
34. Where is the smart person who disagrees — and have I actually heard them out?
35. What's the base rate? What happened to others who made this exact call?
36. Can I run a cheap experiment that produces real information before committing?
37. Five whys deep — do I still have a reason, or just a feeling?
38. What data am I treating as fact that's actually opinion or extrapolation?
39. What would I need to see to change my mind? If nothing would, I'm not testing.
40. Where did this information come from — and what were the incentives of the person who produced it?
41. Is there reason to doubt what I've been told — and have I gone looking for that reason?

05 • FILTER — information hygiene at altitude

You are not deciding with neutral information. Audit the source.

42. Who has financial or career skin in the answer I reach?
43. What direction would each of my key sources want me to lean?
44. Through how many people did this information pass before reaching me?
45. What was probably softened, sharpened, or omitted along the way?
46. What hasn't been brought to me? What's the silence telling me?
47. Who would tell me something inconvenient — and have I actually asked them?
48. When did I last speak directly to a customer, operator, or front-line person on this?
49. Whose framing am I unconsciously adopting? Whose interests does that framing serve?
50. Am I being told what to think — or being given what I need to think for myself?
51. What might someone who strongly disagreed with the conclusion I'm being presented say?
52. What's the one fact, if true, that everyone bringing me information would have an incentive to hide?

06 · PRE-MORTEM — imagining the failure in detail

It is 18 months from now. The decision failed. Why?

53. It's 18 months out and this has failed. Write the autopsy — what does it say?
54. What was the first domino? When could it have been caught?
55. What warning signs did I ignore, rationalize, or fail to instrument?
56. What outside event made it worse — and was it really unforeseeable, or just unconsidered?
57. What did I specifically do, or fail to do, that contributed?
58. Who would I have wanted in the room before deciding, that wasn't?
59. What was the assumption that turned out to be wrong?
60. Was the failure visible early, or only in hindsight? What would make it visible early?
61. What's the failure scenario I'm *not* writing down because it feels too unlikely?
62. For each cause above, what would I put in place *now* to prevent it?

07 · SURFACE — engineer dissent before you decide

Genuine dissent doesn't arise spontaneously. Build the room for it.

63. Have I gathered input from each key stakeholder *privately* before the group meets?
64. Have I spoken last, or did I anchor the room with my view?
65. Who is named — in advance — as the devil's advocate on this decision?
66. Have I asked the specific question "what would make us not do this?"
67. When was the last time someone changed my mind on something important?
68. Who on my team disagrees with me — and do they feel safe saying so out loud?
69. Is the silence in the room agreement, or compliance?
70. Do I have two or three people outside the company who tell me the truth?
71. When I asked for pushback, did I actually *listen*, or just performatively absorb it?
72. If everyone in this room agrees, what's the cost of nobody being wrong?
73. Who in the room is most junior, and have I made it safe for them to disagree with the most senior person?

08 · CONFRONT — what you are really protecting

The honest interrogation. This is where the decision actually is.

74. What am I protecting by making this decision the way I want to?
75. Whose approval am I optimizing for — and would I decide differently if they vanished?
76. What story about myself depends on this going a particular way?
77. What past commitment am I really re-litigating with this decision?
78. Would I decide the same way if I'd never said anything publicly about it?
79. Is fear shaping this? Which fear, specifically: weakness, disappointment, being wrong, irreversibility?
80. Am I deciding from confidence or from urgency? They feel similar from the inside.
81. Is this a decision I keep *not* making? What is my identity protecting me from?
82. What would the version of me who didn't need to be right do here?
83. If I knew this decision wouldn't affect my reputation at all, would I decide differently?
84. In ten years, will I look back on this and wish I'd been braver, or wish I'd been more careful?
85. What would I tell my best friend to do, in this exact situation, if it were theirs?

09 · COMMIT — decide, communicate, calibrate

The decision isn't the moment of choice. It's everything after.

86. Have I written down what I've decided, why, and what I considered?
87. What specific metric or event would tell me this is going wrong?
88. What is my kill criterion — the point at which I would reverse, no matter the sunk cost?
89. When exactly will I review this — and is the date on my calendar?
90. Who needs to hear this from me directly, before anyone else?
91. Have I explained the *reasoning*, not just the conclusion?
92. Have I explicitly said what would change my mind?
93. Am I still selling this decision after it's made? (If so — stop.)
94. Did the people who disagreed feel heard before I committed?
95. When the outcome comes in, will I be able to separate decision quality from luck?

10 · THE FIVE THAT MATTER MOST

If you only ask five questions, ask these. Each one stops at least one common failure mode.

- | | |
|-----|--|
| 96 | Am I solving the right problem, or just the most visible one? |
| 97 | Is this Type 1 or Type 2 — and am I matching my process to that? |
| 98 | Who has skin in this answer, and what would they want me to conclude? |
| 99 | What am I protecting that I don't want to admit I'm protecting? |
| 100 | What would tell me — early — that I got this wrong, and would I actually listen? |

APPENDIX A

The biases that sink CEO decisions

These are not abstract — they are the documented modes by which intelligent, experienced executives consistently make worse decisions than they're capable of. Each one is illustrated by a real story from the history of executive decision-making, drawn from Kahneman, Berger, Munger, and others. Read them before each major decision. Naming them is half the defense.

1. Sunk Cost

Continuing to invest in a failing path because of what's already been spent.

HOW IT SHOWS UP

Polaroid is the cautionary tale. Edwin Land built the company on a single beautiful question — *why do we have to wait for the picture?* — and rode it to dominance for forty years. But as digital photography emerged in the 1990s, Polaroid's leadership kept pouring money into instant film because that was where decades of expertise, factories, and identity were sunk. The question that had built the company was the same one that should have ended it sooner: *why do we still need film at all?* Nobody at the top wanted to ask it. By the time they did, it was 2001 and Polaroid was bankrupt.

THE ANTIDOTE

Antidote: Ask yourself the founder's question — *if I were arriving today with no history, would I invest in this from scratch?* The history is information about how you got here. It is not evidence about whether to stay.

2. Confirmation Bias

Seeking evidence that supports what you already believe; dismissing what doesn't.

HOW IT SHOWS UP

When Brian Chesky and Joe Gebbia pitched Airbnb — strangers paying to sleep on inflatable mattresses in someone's apartment — Paul Graham of Y Combinator, one of the sharpest investors alive, told them flatly: *"no one would want to stay in someone else's bed."* Every fact Graham knew about the hotel industry confirmed his view. The Airbnb founders had been on both sides of the problem themselves — needing a place to stay, having space to rent — and they kept asking the question their evidence base told them to drop. Graham later became one of their seed investors. The lesson cuts both ways: confirmation bias is what made him wrong, and the founders' willingness to test against the consensus is what made them right.

THE ANTIDOTE

Antidote: Assign someone — by name — to argue against you. Spend equal time looking for evidence against your preferred conclusion. Notice *impatience* with disconfirming information as the tell.

3. Expertise Trap

Past expertise becomes the blinders that prevent new questioning.

HOW IT SHOWS UP

Frank Lloyd Wright captured it precisely: *an expert is someone who has stopped thinking, because he knows*. Warren Berger documents the pattern across decades — Polaroid, Kodak, Blockbuster, and a long list of companies whose leaders had genuinely deep knowledge of their industry, and that knowledge was exactly what kept them from seeing what was about to happen to it. The CEO who has "seen this before" is often the one most blind to the case where it's different this time. Reed Hastings was not a video-rental expert when he founded Netflix. The Airbnb founders were not hoteliers. Edwin Land's original Polaroid question came from his five-year-old daughter, not from a camera engineer.

THE ANTIDOTE

Antidote: Periodically ask a question only a naïve outsider would ask. "Why does it have to be this way?" Take it seriously even when — especially when — your expertise has a ready answer.

4. Narrow Framing

Seeing decisions as binary ("should we do X?") instead of among multiple alternatives.

HOW IT SHOWS UP

Edwin Land's instant camera didn't begin with "should we make film faster?" — that was the narrow frame his engineers were trapped in. It began with a different shape of question: *what if the darkroom were inside the camera?* That reframing opened up a problem space the binary couldn't reach. Research on hundreds of executive decisions (Paul Nutt's studies) found that decisions framed as "should we do X, yes or no" failed roughly twice as often as decisions framed with multiple alternatives. The frame is doing more of the work than the analysis.

THE ANTIDOTE

Antidote: Before evaluating, generate at least three genuinely distinct options. If you only have two, you haven't widened the frame — you've narrowed it.

5. Fear of Failure

Optimizing the decision around what protects you from being wrong, rather than what could be right.

HOW IT SHOWS UP

Sebastian Thrun, the engineer behind Google's self-driving car project, told Berger: *"People mainly fail because they fear failure. Innovators have to be fearless."* The fear shows up everywhere in CEO decisions — over-engineering, delaying, hedging, watering down the bold version into a safe version. The single most-cited prompt in Berger's research, asked by everyone from teachers to founders, is the constraint-removing question: *what would you attempt to do if you knew you could not fail?* The point isn't that failure is impossible. The point is that fear of failure is shaping the decision in ways you cannot see while you're inside it.

THE ANTIDOTE

Antidote: Ask the question explicitly: *what would I do here if I knew I couldn't fail?* The gap between that answer and your current answer is the size of fear's influence on this decision.

6. Premature Answer

Jumping to the solution before the question is well-formed.

HOW IT SHOWS UP

Sakichi Toyoda built Toyota on a discipline so simple it sounds childish: when something goes wrong, ask Why five times. A faulty part comes off the line. Why? An operator made a mistake. Why? Insufficient training. Why? Underfunded program. Why? Misallocated budget. Why? Wrong company priorities. The first answer is almost never the real one. Eric Ries, who built the Lean Startup methodology around the same practice, told Berger it works *"because it's designed to overcome the limits of human psychology — we tend to personalize systemic problems."* The first answer blames a person; the fifth answer reveals a system. Most CEO decisions are made at the level of the first answer.

THE ANTIDOTE

Antidote: Before committing, ask Why three to five times about your own conclusion. "Why is this the answer?" — and then "why is *that*?" — until you reach something structural, not personal.

7. Gut-Trust Bias

Treating gut feeling as analysis, especially when stakes are high.

HOW IT SHOWS UP

Ron Shaich, the founder and former CEO of Panera Bread, told Berger that a CEO's *"primary occupation must be to discover the future"* — and that questioning, not instinct, is the only tool for navigating uncharted terrain. He was making a specific point: the old model where the senior leader walked in with confidence, experience, and gut was useful in a stable environment. It is dangerous in what he called the VUCA world — volatile, uncertain, complex, ambiguous. In conditions where the rules are changing, gut is mostly a record of how things used to work. Berger cites current research: trusting gut feelings is consistently worse than deliberate analysis when stakes are high and the environment is changing.

THE ANTIDOTE

Antidote: When you feel certain, slow down. Treat the gut signal as a hypothesis, not a conclusion. Test it against an explicit alternative before acting.

8. Groupthink

Convergence toward consensus, especially when the leader has signaled a preference.

HOW IT SHOWS UP

Berger documents the pattern repeatedly: companies start as questioners — every great one in the book did — and then a hierarchy forms, a methodology is established, rules are set, and the question that built the company becomes the question no one is allowed to ask anymore. *"Companies are like people,"* he writes. *"They start out doing it, then gradually do it less and less."* The mechanism is social, not analytical. Disagreeing with the CEO is costly. Agreeing is free. So the room converges, and the convergence is mistaken for alignment when it is actually compliance. The biggest decisions get made in rooms where no one disagrees because no one can afford to.

THE ANTIDOTE

Antidote: Solicit views privately before the group meets. Speak last, not first. Visibly reward the person who disagreed — once. Word travels fast in both directions.

9. Optimism Bias

Systematically overestimating the likelihood of success and underestimating time, cost, and obstacles.

HOW IT SHOWS UP

Kahneman calls it the planning fallacy — and it is the most pervasive bias in executive decision-making. Studies of large infrastructure projects found that nine out of ten came in over budget. Studies of corporate mergers found that 70–90% destroy shareholder value — yet every acquiring CEO's financial model showed the deal working. The people making these decisions were not stupid. They were human. The inside view of their own situation was too vivid and too optimistic, and the outside view — what actually happens to people who do this — was never seriously consulted. Theranos is the extreme version: Elizabeth Holmes's overconfidence in her own vision led her to suppress accurate test results rather than confront evidence that the technology did not work. The pattern is not unique to fraudsters. It shows up in every ambitious project plan ever written.

THE ANTIDOTE

Antidote: Force the outside view. Before committing, ask: what is the base rate for projects like this? What percentage of comparable decisions actually worked out? Take that number seriously even when your situation feels different — it always feels different from the inside.

10. Overconfidence

Systematic overestimation of your own knowledge, judgment, and the precision of your forecasts.

HOW IT SHOWS UP

Kahneman asked a group of experts to give 98% confidence intervals for factual questions — ranges they were 98% certain contained the right answer. The true answers fell outside those intervals roughly 45% of the time. That is not a rounding error; that is a fundamental miscalibration. The same pattern holds for CEOs forecasting revenue, lawyers forecasting trial outcomes, and engineers estimating project timelines. The fix is not modesty — it is *calibration*. Philip Tetlock's decades of forecasting research found that a small group of people are genuinely better calibrated than domain experts — not because they know more, but because they track their predictions, update when wrong, and express uncertainty in numbers rather than words like 'very likely'. The difference between a well-calibrated and a poorly-calibrated decision-maker is not knowledge. It is the habit of checking.

THE ANTIDOTE

Antidote: Replace words with numbers. 'Very likely' means nothing. '75%' forces honesty. Track your probability estimates against outcomes over time — the act of tracking is itself what builds calibration.

11. Framing Effect

The same information, presented differently, produces different decisions.

HOW IT SHOWS UP

A medical procedure with a 90% survival rate sounds more appealing than one with a 10% mortality rate. Mathematically identical. Psychologically opposite. Kahneman and Tversky documented this in the 1980s and it remains one of the most replicated findings in behavioral economics. For a CEO: the way a recommendation is framed — gains versus losses, percentages versus absolute numbers, "80% chance of success" versus "one in five chance of failure" — determines the decision as much as the underlying data. The people briefing you have almost certainly framed the information in the direction that produces the answer they want. This is not always malicious; it's often unconscious. But the effect on your decision is the same either way.

THE ANTIDOTE

Antidote: Restate every important number in its opposite framing before deciding. If a proposal is presented as upside, force yourself to articulate the downside in the same terms — and vice versa. Ask: how would this recommendation read if the person presenting it wanted you to say no?

12. Inconsistency Avoidance

Once a position is taken publicly, people resist changing it — constructing elaborate justifications rather than updating.

HOW IT SHOWS UP

Charlie Munger called this one of the most underrated human tendencies: once we commit to a belief or a decision — especially publicly — we become its defenders rather than its evaluators. The psychological cost of changing our minds is experienced as a loss of identity, not as a correction. The result: executives who privately know a strategy is failing will continue funding it, framing each failure as a temporary setback, because reversing course means admitting the original call was wrong. The organizational cost of one person's inconsistency avoidance can be enormous. It is distinct from sunk cost (which is about money already spent) — this is about ego and identity already invested. The question 'what would I do if I were arriving fresh today?' is the cleanest diagnostic because it removes the identity stake from the decision.

THE ANTIDOTE

Antidote: Separate the quality of the original decision from the quality of the current decision. Changing course is not failure — it is new information properly weighted. Create a culture where updating publicly is rewarded, not penalized. The leader who says 'I was wrong about X, here is what I now know' builds more trust than the one who quietly buries it.

APPENDIX B

The decision journal

Annie Duke's recommendation, and the single most powerful tool for getting better at decisions over time. Before committing, write down — in plain language — the answers to the following. Dated, sealed, and unread until your review date.

- ◆ **The decision.** What am I deciding, in one sentence?
- ◆ **The reasoning.** Why this option, and not the others I considered?
- ◆ **The expectation.** What do I expect to happen, and on what timeline?
- ◆ **The probability.** What's my numerical estimate of success? (Force a number. "Pretty likely" doesn't calibrate.)
- ◆ **The disconfirmers.** What are the top three reasons I might be wrong?
- ◆ **The emotional state.** What state am I in as I make this? Tired, rushed, frustrated, euphoric, fearful? Name it. State distorts.
- ◆ **The information environment.** Whose input shaped this most? Whose perspective am I missing or under-weighting?
- ◆ **The success criteria.** What would have to be true at the review date for me to call this decision well-made — regardless of outcome?
- ◆ **The kill criterion.** What would tell me this needs to be reversed or abandoned?
- ◆ **The review date.** When will I revisit this entry — and who else will I show it to?

WHY THIS MATTERS MORE THAN IT SOUNDS

Without a written record of what you believed when you decided, you cannot tell — months later — whether you actually decided well. Memory edits the past to match outcomes. The journal is the only honest mirror. Six entries in, you start seeing your patterns. A year in, you become a measurably better decision-maker. There is no faster way.

APPENDIX C

When not to use this

Process is not free. Applied to the wrong decisions, Before You Commit will slow you down without making you better. Skip or compress it when:

- 1. The decision is genuinely Type 2 and easily reversible.** Spending three days on a decision you can unwind in a week is a tax on your company. Move.
- 2. The decision is below your altitude.** If a direct report should own this, the right "decision" is delegation. The most expensive CEO time is time spent on decisions that aren't yours.
- 3. Speed itself is the strategic edge.** Some decisions degrade rapidly with deliberation — competitive responses, crises, time-bound opportunities. Apply judgment, not full process.
- 4. You're stalling.** If you've already done the work and are reaching for a seventh framework, you're not improving the decision — you're avoiding it. The right move is to commit.
- 5. It's actually a values question, not a decision.** Some "decisions" are really questions about who you want to be. Frameworks won't help. Sit with the question instead.
- 6. You don't own the decision.** If the choice genuinely belongs to someone else — a board, a co-founder, a regulator — and you are providing input rather than deciding, use this guide to sharpen your input, not to drive the outcome. Trying to run a full decision process when you don't hold the D creates confusion and erodes trust.
- 7. The relationship is the decision.** When the real question is whether to trust someone — a new hire, a co-founder, a partner — no analytical framework substitutes for direct, extended observation of how they behave under pressure. Use the guide to structure what you're looking for, then put the framework down and watch.

APPENDIX D

Self-calibration

Without measurement, you can't tell if this process is actually improving your decisions or just making you feel more thorough. The discipline below is what separates compounding decision-makers from ones who plateau.

The quarterly review

Once a quarter, pull out your last several decision journal entries and ask:

1. Of the decisions whose review dates have passed, how many of my probability estimates were in the right zone? Was I systematically over- or under-confident?
 2. Where did my expected timeline diverge most from what actually happened? What does that tell me about how I forecast?
 3. Which of my "top three reasons I might be wrong" actually showed up? Was I anticipating real risks — or comfortable ones?
 4. On the decisions that went badly: was it a bad decision, or good decision and bad luck? Be honest.
 5. On the decisions that went well: was it a good decision, or bad decision and good luck? Even more honest.
 6. What's the single pattern across my decisions that I keep getting wrong? That's the area to focus on next quarter.
-

APPENDIX E

Question techniques from the literature

A curated set of question patterns documented in Warren Berger's *A More Beautiful Question*. These are not additional items in the main 100 — they are **techniques**: ways of asking that change the kind of answer you get. Use them when the standard questions aren't unlocking what you need.

PART I — FRAMING, QUESTIONING & OPTION GENERATION

How Might We...?

Min Basadur (Procter & Gamble, 1970s) → IDEO → Google → Facebook

Three words, each doing specific work. **How** assumes there are solutions — it creates creative confidence. **Might** defers judgment, so people propose ideas without committing to them. **We** says we'll figure this out together. Replace "how can we" or "how should we" — both of which imply judgment — with "how might we." Tim Brown of IDEO uses this on every project. It is the single most broadly adopted questioning technique in modern business.

EXAMPLES

- › How can we beat the competitor? → **How might we** serve a customer they're ignoring?
- › How should we cut costs? → **How might we** deliver the same value with fewer inputs?
- › How do we fix retention? → **How might we** make this product worth coming back to?

Why → What If → How

Edwin Land's pattern at Polaroid, generalized by Berger as the inquiry sequence

Most decisions get attacked at the How stage — execution — when the breakthrough is hiding upstream. Berger's documented pattern across innovation cases: start with **Why** (challenge the assumption: why does it have to be this way?), move to **What If** (generate alternatives without committing), then arrive at **How** (translate the chosen direction into action). Skipping straight to How is the most common mistake — and the most expensive.

EXAMPLES

- › **Why** are we losing this market? → **What if** the customer isn't who we think it is? → **How** would we serve the real customer?
- › **Why** can't we ship faster? → **What if** we ship something smaller, more often? → **How** would we restructure to do that?

Hero–Dragon–Quest (The Problem Frame)

Arnaud Chevallier & Albrecht Enders · Solvable: A Simple Framework for Complex Problems, 2022

A well-formed problem frame has exactly four components. The **Hero**: who is making this decision, and what authority do they actually hold? The **Treasure**: what specific outcome are they trying to achieve — stated as an end state, not a solution? The **Dragon**: what is the specific obstacle standing between the current state and the treasure? And the **Quest**: the overarching question that must be answered to slay the dragon and reach the treasure. Chevallier and Enders found that executives almost always skip the Dragon. They jump from Treasure (the outcome they want) directly to Quest (the question they think they're answering) — without ever naming what is actually in the way. The result: solutions that address a comfortable version of the problem rather than the real constraint. You cannot design a good solution for a dragon you haven't named. The frame is also a diagnostic for whether the problem is worth solving: if you can't state the dragon clearly, you don't yet understand the problem well enough to solve it. **Worked example — a logistics company's growth problem.** The CEO walks in with a stated problem: 'We need to grow revenue by 30%.' The team immediately generates solutions: new sales hires, new markets, price increases. None of them work. Applying the Hero-Dragon-Quest frame: *Hero*: The CEO, with full P&L authority and a board mandate to grow. *Treasure*: Sustainable 30% revenue growth within 24 months without proportional cost increases — i.e., profitable growth, not just top-line. *Dragon (unnamed until now)*: Customer churn among the top 20 accounts is running at 35% annually — meaning the company is replacing nearly a third of its best revenue every year just to stay flat. Every new sale is filling a leaking bucket. New sales hires accelerate the inflow but do nothing about the leak. The dragon isn't market access or sales capacity. It is retention. *Quest*: Why are the top 20 accounts churning, and what would make them stay and expand? The original framing ('how do we grow?') generated solutions that addressed the treasure while ignoring the dragon. The reframed quest ('why are we losing our best customers?') generates a completely different investigation — and a completely different answer. The company discovers that churn is driven by a single integration failure in their platform that only affects accounts above a certain volume threshold. Fixing it costs less than one new sales hire and unlocks the growth that no amount of new sales could have achieved.

EXAMPLES

- › *Who is the Hero — and do they actually hold the authority this decision requires?*
- › *What is the Treasure — stated as an outcome, not a solution? Strip any embedded solutions from the statement.*
- › *What is the Dragon — the specific obstacle between current state and treasure? If you can't name it precisely, you're not ready to solve.*
- › *What is the Quest — the single question that, if answered, would show you the path past the Dragon?*
- › *Are you designing solutions for the Dragon you named — or for a more comfortable, less threatening version of it?*

Consider the Opposite

Richard Larrick (decision research) — adopted by Berger as a bias-awareness move

When you feel strongly about a position, force yourself — briefly — to argue the opposite. Not to change your mind, but to neutralize the lean. The technique works by activating a part of your reasoning that confirmation bias normally suppresses. Larrick's research shows this single move reduces biased judgment more than almost any other intervention. Berger captures it as a question template: *knowing that I lean one way, what would "opposite me" think?*

EXAMPLES

- › *I'm certain we should acquire this company. → What would Opposite Me say about why this is a mistake?*
- › *This hire is obvious. → What would a skeptic see that I'm missing because I want this to work?*

Questions to Stop Asking (the toxic ones)

Berger's leadership inquiry section: questions that look productive but kill thinking

Some questions sound rigorous and feel responsible to ask, but they shut down the conversation rather than open it. Berger flags three patterns common at the top of organizations — and the better version of each. The reframe is small. The downstream effect on what your team brings you is enormous.

EXAMPLES

- › *What's wrong with this? → What are we doing well, and how might we build on that?*
- › *Whose fault is it? → How can we work together to shore up the weak link?*
- › *Haven't we tried this already? → If we tried this now, what might be different — and how might that change the result?*

The Constraint-Removal Question

Question Berger documents across founders, athletes, scientists, and creatives

What would I attempt to do if I knew I could not fail? The point is not that failure is impossible — the point is that fear of failure is shaping the decision in ways you cannot see while you are inside it. The question removes that constraint temporarily. The gap between that answer and your current answer is the size of fear's influence on this decision. Sebastian Thrun put it to Berger directly: *people mainly fail because they fear failure.*

EXAMPLES

- › *What would I do if I knew I could not fail?*
- › *What would I do if money were no object?*
- › *What would I do if I had only one dollar to spend?*
- › *What would I do if I had to ship this in 48 hours?*

Stepping Back

Berger's account of why distance produces clarity — the Edwin Land walking principle

Land's instant-camera breakthrough did not happen at a desk. It happened on a long walk, alone, after his daughter asked the original question. Berger documents the same pattern across innovators: the answer arrives in the shower, on a run, on a walk, in the moment of detachment. The mechanism is not magic — it is that stepping back disables the part of the mind that's been defending the existing framing, and lets a different framing surface. For a CEO, the operational version of this is: **do not commit to a major decision in the meeting where it is discussed**. Leave the room. Walk. Decide later.

EXAMPLES

- › *If I left this meeting and walked for an hour, what would I see that I'm not seeing now?*
- › *What would I decide if I had to wait 24 hours before saying yes or no?*
- › *Am I in the right state to decide this — or in the state that brought the urgency into the room?*

Chesterton's Fence

G.K. Chesterton (1929), operationalized by Gefroh and others as a decision discipline

Before removing any rule, constraint, process, or structure — understand why it exists. Chesterton's original formulation: the reformer who sees a fence across a road and says '*I don't see the use of this; let us clear it away*' should be stopped. The correct response is: '*If you don't see the use of it, I certainly won't let you clear it away. Go away and think. When you can come back and tell me that you do see the use of it, I may allow you to destroy it.*' For a CEO: if you can't explain the function a thing serves, you're not yet ready to change it. This applies equally to org structures, approval processes, team configurations, and product constraints that 'everyone agrees' are unnecessary.

EXAMPLES

- › *Why does this constraint/rule/structure exist? Who put it there, and what were they trying to prevent?*
- › *If I remove this, what happens the first time the thing it was protecting against occurs?*
- › *Am I solving a real problem — or removing friction that was load-bearing?*

Multi-Frame Diagnosis

Gefroh's framing-as-lens concept — the mechanism behind most misframed decisions

Your default frame — people, process, technology, culture, market, systems — determines which causes you can see and which solutions you can reach for. Most executives use one or two frames without knowing it. The CRO framing a churn problem as an acquisition problem. The engineering manager framing a planning problem as a PR-review problem. The mistake is invisible from inside the frame. **The antidote is deliberate multi-frame diagnosis:** take the same situation and run it through at least three different lenses before settling on a cause. The solution you couldn't see before is usually in the frame you haven't tried. **A worked example — the same problem, four diagnoses.** A B2B SaaS company is missing its revenue targets for the third consecutive quarter. The leadership team frames it differently depending on who is in the room. *People frame (default for most CEOs):* the sales team is underperforming. Diagnosis: wrong hires, poor coaching, low accountability. Solution: replace the VP of Sales, rebuild the team. Cost: 12 months and significant disruption — and if the frame is wrong, the new VP inherits the same unresolved problem. *Process frame:* the sales cycle is too long and deal velocity has dropped 40% in 18 months. Diagnosis: too many approval steps, unclear pricing authority, proposals stuck in legal review. Solution: streamline the close process, give AEs more pricing discretion. Faster to implement, lower risk. *Market frame:* three well-funded competitors entered the segment last year and are discounting aggressively to buy market share. Diagnosis: positioning erosion — the company is no longer the obvious choice in a crowded field. Solution: sharpen ICP, exit the contested mid-market, double down on a defensible niche. This reframes the entire go-to-market, not just the sales motion. *Product frame:* feature adoption data shows that customers who use the core workflow daily renew at 94%; those who don't churn at 61%. The problem isn't selling — it's that the product isn't delivering the promised value fast enough. Solution: redesign onboarding to drive activation in week one. This frame makes the revenue problem a product problem — invisible from the sales frame. All four diagnoses are partially correct. Only one or two will be load-bearing. The CEO who runs only the people frame fires the VP of Sales and solves nothing. The CEO who runs all four picks the highest-leverage intervention — or sequences them correctly.

EXAMPLES

- › *What is my default frame for this problem? Name it explicitly before diagnosing.*
- › *Through a people lens: what human behavior, motivation, or capability is causing this?*
- › *Through a systems lens: what structural incentive or feedback loop is producing this outcome?*
- › *Through a market lens: what external pressure or change is the actual cause?*
- › *After running all three: which frame gives me the most actionable diagnosis?*

Examine Bright Spots

Thomas Wedell-Wedellsborg, What's Your Problem?, 2020 · also: Chip & Dan Heath, Switch, 2010

When diagnosing a problem, most people ask: *where is it going wrong?* Bright spots flips that. Ask instead: *where is this problem **not** occurring?* What's working — even partially — and why? The logic: if the problem were universal and systemic, there would be no bright spots. The fact that they exist tells you something the problem itself cannot: that a solution is already operating somewhere in your organization. You don't need to invent it — you need to find it, understand it, and scale it. BarkBox reframed a shelter-dog adoption problem from an advertising problem to an access problem by asking where adoptions were working — and found that mobile access, not awareness, was the real constraint. The result was a dating-app-style adoption app built for \$8,000 that achieved outsized impact. Also: negativity bias (the scientific tendency to focus on what's failing) makes bright spots systematically invisible. You have to look for them deliberately.

EXAMPLES

- › *Where is this problem not occurring — and what is that place doing differently?*
- › *Has this problem ever been solved, even partially, somewhere in this organization? What can we learn from that?*
- › *Which person, team, or market is performing well despite having the same constraints as everyone else?*
- › *If I focused entirely on replicating the bright spot rather than fixing the problem, what would I do?*

Five Reframing Moves

Thomas Wedell-Wedellsborg, What's Your Problem?, Harvard Business Review Press, 2020

When a problem feels stuck — when you've analyzed it correctly and still can't find the right answer — the issue is usually the frame, not the analysis. Wedell-Wedellsborg's five moves are a rapid diagnostic for challenging the frame you've inherited. Use them in sequence or in parallel when Stage 01's framing diagnostic hasn't cracked it open.

EXAMPLES

- › **Look outside the frame.** *Is the stated problem actually true? Is information missing? Are you imposing a constraint that isn't real? Ask: how do we know this is the problem?*
- › **Rethink the goal.** *Is the stated objective actually the right one? Intel tried to make Hawking's text system faster — he asked them to make it slower, so he could maintain eye contact while composing. More of an obviously good thing was the wrong direction entirely.*
- › **Examine bright spots.** *Where is the problem not occurring? What's working — and why?*
- › **Look in the mirror.** *What role are we playing in creating or sustaining this problem? The factory manager who opened every exit-interview with 'tell me how the company failed you' got better information than anyone looking for someone to blame.*
- › **Take their perspective.** *What does this situation look like from the other person's position — not just emotionally, but contextually? What do they know and not know? What are they actually trying to get done?*

RAPID — Who Has the D?

Paul Rogers & Marcia Blenko (Bain & Company) · via McGrath, The Little Book of Big Decision Models, 2016

The single most common reason group decisions fail is not bad analysis — it's ambiguity about who is actually deciding. RAPID makes decision roles explicit and assigns them to named individuals *before* the decision is made. The five roles: **Recommend** (develops the proposal), **Agree** (must sign off — veto power), **Perform** (executes once decided), **Input** (consulted but no veto), **Decide** (the D — one person, final call). The CEO's most important use of RAPID: resist defaulting to being the D on every decision. When the CEO is always the D, the organization stops deciding. Every senior decision should have an explicit owner who is not you — unless the stakes genuinely require it.

EXAMPLES

- › *Who is the D on this decision? If you can't name one person, you don't have a decision — you have a discussion that will drift.*
- › *Who has veto (Agree) — and have they been consulted before the recommendation is formed, not after?*
- › *Who is Input-only — and do they know they're Input-only, or are they expecting a vote they won't get?*
- › *Should I be the D here — or am I taking the D because I haven't trusted someone enough yet to own it?*

De Bono's Six Thinking Hats

Edward de Bono (1985) — the most widely used parallel-thinking framework in executive practice

Six Hats forces a group to think about a decision from six explicitly different modes, sequentially — preventing the mixed, unproductive debate where everyone argues from a different position simultaneously. Each hat is a color; each color is a mode. You wear one hat at a time, as a group, and move through them deliberately. The result: richer diagnosis, fewer ego battles, and no idea goes unexamined because it came from the wrong person. Most powerful for team decisions where advocacy and emotion are already crowding out clear thinking.

EXAMPLES

- › **White Hat — Data:** *What do we know? What information is missing? No interpretation — facts only.*
- › **Red Hat — Emotion:** *What does your gut say? What feelings and intuitions are in the room? No justification required.*
- › **Black Hat — Caution:** *What could go wrong? What are the risks, weaknesses, and dangers of this decision?*
- › **Yellow Hat — Optimism:** *What's the best-case scenario? What value does this create if it works?*
- › **Green Hat — Creativity:** *What else might we do? What alternatives haven't been considered?*
- › **Blue Hat — Process:** *What thinking do we need? Who facilitates? What's the next step? (Usually worn by the chair.)*

Socratic Questioning

Socrates (469–399 BC) — five-category taxonomy for structured inquiry in high-stakes proposals

Socratic questioning probes assumptions, tests evidence, and exposes inconsistencies without the questioner asserting a position. For a CEO, it's the discipline of asking the questions that reveal what's being taken for granted — especially in a room where people are invested in a conclusion. Apply the five categories in sequence when examining any high-stakes recommendation before committing. The sequence matters: clarify first, then probe assumptions, then evidence, then perspectives, then consequences.

EXAMPLES

- › **1 — Clarification:** *What do you mean exactly? Can you give a concrete example? Can you say that another way? What is the central issue? What are we actually deciding here?*
- › **2 — Probing Assumptions:** *What are you assuming? What would happen if that assumption were false? What are we taking for granted that hasn't been examined? How did you reach that conclusion? Is there another way to look at this that you haven't tried?*
- › **3 — Probing Evidence:** *Why is that happening? What evidence supports this claim? What evidence would contradict it — and have we looked? Is there reason to doubt this? Where did this information come from, and who produced it? What other information do we need before we can be confident?*
- › **4 — Questioning Perspectives:** *What might someone who strongly disagreed with this say? What would a competitor, customer, or regulator see that we're missing? Have we always thought this way — or did something change? What would we think if this came from a different source?*
- › **5 — Implications & Consequences:** *Then what would happen? What are the second-order consequences of this decision? How does this fit with what we already know and have committed to? Who else is affected by this — and have they been heard? Could this really happen, or are we confusing plausible with probable?*
- › **6 — Questioning the Question:** *Why do you think I asked that? Am I asking the right question here, or is there a better one? What question should we actually be asking that we haven't yet? How does this question link to the outcome we're really trying to achieve? (Use this category at the end of any discussion to test whether the inquiry itself was well-aimed.)*
- › **Using it in a room:** *Ask one question at a time. Wait for a full answer before asking the next. The power is in the pause — what people say when they feel genuinely heard is usually more accurate than what they lead with.*

Ladder of Inference

Chris Argyris (1970s) · popularized by Peter Senge, The Fifth Discipline, 1990

The Ladder of Inference describes the mental pathway from raw data to action — and why most disagreements happen not at the level of facts but at the rungs above them. The seven rungs: (1) **Raw data** — what actually happened. (2) **Selected data** — what you chose to notice (ignoring the rest). (3) **Interpreted data** — what you think it means. (4) **Assumptions** — what you add from your existing beliefs. (5) **Conclusions** — what you decide is true. (6) **Beliefs** — what this reinforces about the world. (7) **Actions** — what you do. Most executives operate at rung 7 while believing they are responding to rung 1. The diagnostic power: when two people reach different conclusions from the same meeting, they diverged somewhere on the ladder — and the disagreement cannot be resolved by restating their conclusions louder. You have to find the rung where they split.

EXAMPLES

- › *What is the raw data here — what actually happened, before any interpretation?*
- › *What data have I selected to focus on — and what am I not giving equal weight to?*
- › *What am I assuming that I haven't made explicit, and why do I believe it?*
- › *Is my conclusion sound — or have I climbed the ladder so fast I've lost sight of the original data?*
- › *Where on the ladder did I and the person I disagree with diverge? That's the real conversation.*
- › *What are the actual facts — separating what was observed from what was inferred?*

Regret Minimization Framework

Jeff Bezos, 1994 — used personally when deciding to leave D.E. Shaw to start Amazon

Project yourself to age 80 and look back. Which choice would you regret more — having tried and failed, or never having tried? Bezos used this explicitly: at 80, he knew he wouldn't regret not having participated in the internet gold rush if he had stayed. But he would regret it if he hadn't tried. The decision became easy. The framework is distinct from 10-10-10 (which manages short-term emotional pressure) — Regret Minimization operates on a longer horizon and speaks specifically to identity and legacy rather than comfort. It is most useful for decisions where the status quo feels safe but may be a form of avoidance: not taking the risk, not making the change, not having the conversation. The question cuts through the daily noise and asks what the decision looks like from outside the anxiety of the present moment.

EXAMPLES

- › *At age 80, looking back — which choice would I regret more?*
- › *Am I choosing the status quo because it's genuinely right, or because risk feels uncomfortable from the inside?*
- › *What would I think of this decision if I were watching myself make it from a distance of 40 years?*
- › *Is the fear I'm feeling today the kind that will seem important at 80 — or the kind that vanishes as soon as the decision is made?*

10/10/10 — Temporal Perspective

Suzy Welch, 10-10-10: A Fast and Powerful Way to Get Unstuck in Love, at Work, and with Your Family, 2009

How will you feel about this decision in 10 minutes? In 10 months? In 10 years? The three time horizons expose the distortion created by present-bias — the systematic tendency to overweight immediate discomfort or reward relative to future consequences. A decision that feels impossible in 10 minutes often looks obviously correct in 10 years. A decision that feels obviously correct in 10 minutes often looks catastrophic in 10 months. The technique is distinct from Regret Minimization: where Regret Minimization operates on identity and legacy, 10/10/10 is a rapid calibration tool for cutting through the emotional noise of the present moment — especially useful for decisions where urgency or social pressure is doing most of the work. The 10-minute horizon names the feeling driving you right now. The 10-month horizon names the operational consequence. The 10-year horizon names what actually matters. When all three align, the decision is clear. When they diverge — particularly when 10 minutes says one thing and 10 years says the opposite — you are almost certainly looking at a present-bias distortion rather than a genuine strategic signal.

EXAMPLES

- › *In 10 minutes, how will I feel about this decision — and is that feeling useful information, or just the emotion of the moment?*
- › *In 10 months, what will the operational consequences of this decision be? What will I be dealing with that I'm not thinking about now?*
- › *In 10 years, will this decision matter? If not — move fast. If yes — slow down precisely here.*
- › *Where do the three horizons diverge? The gap between 10-minute anxiety and 10-year clarity is the size of present-bias in this decision.*
- › *If I could fast-forward to 10 months and look back on this moment, what would I wish I had done?*

Optionality — Does This Decision Open or Close Doors?

Nassim Taleb (Antifragile, 2012) · Real options theory · Chess positional strategy

Every decision does two things: it produces an outcome, and it changes the set of decisions available to you in the future. Most decision frameworks focus entirely on the first. The chess master focuses as much on the second: a move that wins material but locks your position into a single forced line is often weaker than a move that wins less but keeps multiple paths open. The principle generalizes directly to business. **A decision that preserves optionality — that keeps future doors open — has value beyond its immediate outcome.** A decision that forecloses options has a cost beyond its immediate price. This is especially important under uncertainty: when you don't know which path will be right, the decision that keeps more paths available is usually better than the one that commits you fully to the path that looks best now. Taleb's formulation: seek situations where the downside is capped and the upside is open-ended. Avoid situations where the upside is capped and the downside is open-ended. The practical executive version: before committing, ask whether you are buying an option or selling one. Acquisitions, exclusive contracts, and long-term bets often sell optionality — locking you into a path in exchange for certainty. Pilots, experiments, and partnerships often buy it. Neither is always right. But knowing which you are doing is non-negotiable.

EXAMPLES

- › *Does this decision open future paths or close them? Which doors does it lock?*
- › *Under the uncertainty I'm facing, is buying optionality more valuable than committing? What would it cost to keep the option open a little longer?*
- › *Am I selling optionality without pricing it correctly — committing to a path whose constraints I haven't fully mapped?*
- › *What decisions will this make easier in 12 months? What decisions will it make harder or impossible?*
- › *If circumstances change significantly in the next year, is this decision reversible — or will I be locked in?*
- › *Is there a version of this decision that achieves 80% of the upside while preserving significantly more optionality?*

PART II — TESTING, RISK & DECISION QUALITY

Weighted Criteria — Decide Before You Evaluate

McKinsey / Barbara Minto tradition · Chevallier & Enders, Solvable, 2022 · anchoring research

The most common mistake in option evaluation is applying criteria *after* you already know which option you prefer. At that point, you unconsciously adjust the weights to match your preference — this is anchoring, and it is nearly impossible to resist in real time. The antidote is structurally simple but psychologically difficult: **define your criteria and their relative weights before you evaluate the options**. Force yourself to answer: what would a good decision look like, independently of which option delivers it? Typical criteria for executive decisions: strategic fit, implementation feasibility, financial return, time to impact, reversibility, team capability, and stakeholder alignment. The weights reveal your actual priorities — not your stated ones. Two additional disciplines: (1) separate the weighting session from the evaluation session by at least a day, so the options don't contaminate the criteria. (2) if two people weight the criteria differently, the disagreement about weights is more important than the disagreement about options — it means they're optimizing for different things. Resolve the weighting first.

EXAMPLES

- › *What are my criteria for a good decision here — stated without knowing which option scores highest on each?*
- › *How do I weight them relative to each other? If I could only optimize for one, which would it be? That's the most important criterion.*
- › *Have I set the weights before examining the options — or am I working backward from a preferred answer?*
- › *Where do my stated weights diverge from my real weights? The real weights are revealed by which trade-offs actually feel acceptable.*
- › *If two people in the room weight the criteria differently, what does that tell us about what we're actually optimizing for?*
- › *After scoring options against weighted criteria: if the 'winner' surprises me, do I update my preference — or update the weights? Be honest.*

MECE Thinking

McKinsey methodology — Mutually Exclusive, Collectively Exhaustive — formalized by Chevallier & Enders in Solvable

MECE (pronounced 'me-see') is the discipline of organizing options or causes so that they don't overlap (mutually exclusive) and together cover every possibility (collectively exhaustive). Mutual exclusivity forces clarity about how alternatives relate to each other. Collective exhaustivity prevents blind spots — ensuring you haven't forgotten a whole category. Applied to CEO decisions: before evaluating options, check that they are genuinely distinct (not three flavors of the same move) and that together they cover the real option space (not just the ones that came to mind first). Most executive option-sets fail the CE test — they're exhaustive-looking but have large gaps.

EXAMPLES

- › *Are these options genuinely distinct — or am I evaluating three versions of the same move?*
- › *Have I covered the full option space, or just the options that came to mind first?*
- › *What's the category of option I haven't named yet? (The gap in CE is usually where the best answer lives.)*
- › *If I drew a map of all possible causes, would my diagnosis account for all of them?*

LEAD — Testing Your Hypotheses

Chevallier & Enders (Solvable, 2022) — the four-step evidence protocol for diagnosis

Once you've mapped the possible causes of a problem (the why-map), each candidate cause is a hypothesis, not a fact. LEAD is the discipline for testing them. The crucial instruction: **favour opposing evidence**. Focusing on supporting evidence promotes overconfidence. Your job is to find the evidence that would change your mind — and take it seriously when you do. Note: even proving a hypothesis true doesn't exonerate others. Problems often have multiple real causes simultaneously.

EXAMPLES

- › **L — Locate:** What evidence is relevant to this hypothesis? Identify it before looking — avoid cherry-picking.
- › **E — Evaluate:** What is the quality of each piece of evidence? Who produced it? What incentives shaped it?
- › **A — Synthesize:** What does the body of evidence say as a whole? Identify where it supports and where it opposes the hypothesis.
- › **D — Decide:** Accept, reject, or require more evidence. State explicitly what would change your mind.

Backcasting (the Parade)

Heath & Heath (Decisive, 2013) · Annie Duke (How to Decide, 2020)

The pre-mortem imagines failure and works backward. Backcasting is the companion move: imagine the decision has *succeeded* — and work backward from that. The two techniques together give a complete view of the future: you need both the failure autopsy and the success parade to avoid asymmetric planning. Backcasting is also useful as a motivational planning tool — the reasons you'll succeed are the things to build in advance.

EXAMPLES

- › *It is 18 months from now. This decision has worked brilliantly. What did you do that made it succeed?*
- › *Looking back from that future: which of your own decisions and actions mattered most?*
- › *What external conditions helped? Which of those can you influence now?*
- › *What does the backcast tell you to build in, pre-commit to, or protect — that the pre-mortem alone wouldn't reveal?*

Resulting — Don't Judge Decisions by Outcomes

Annie Duke (Thinking in Bets, 2018 · How to Decide, 2020)

Resulting is the tendency to use outcome quality to judge decision quality. You run a red light and get through safely — great decision? You go through a green and get hit — terrible decision? No. Resulting causes you to learn the wrong lessons from experience: repeating bad-process-lucky-outcome decisions, abandoning good-process-unlucky-outcome ones. The antidote is separating the two explicitly, both when evaluating your own decisions and when evaluating others'. A good decision can produce a bad outcome. A bad decision can produce a good outcome. Judging yourself and your team on outcomes alone systematically corrupts the feedback loop.

EXAMPLES

- › *Was this a good decision or a good outcome? These are different questions.*
- › *What was the quality of the process at the time — based only on what was known then?*
- › *Am I updating my process based on this result, or updating my assessment of luck?*
- › *If this decision were made 100 times, what percentage would produce this outcome?*

Second-Order Thinking


Shane Parrish (Farnam Street) · Howard Marks (The Most Important Thing) · Ray Dalio (Principles)

Most people stop at first-order consequences — the immediate, obvious effect of a decision.

Second-order thinking asks: and then what? Third-order: and after that? Many decisions that look good at the first order create serious damage at the second. Many decisions that look bad at the first order are good at the second. Ray Dalio: *'Failing to consider second- and third-order consequences is the cause of a lot of painfully bad decisions, and it is especially deadly when the first inferior option confirms your own biases.'* For CEO decisions: always map at least two orders before committing. The second order is where competitive response, regulatory reaction, cultural backlash, and system-level effects live. **A worked example.** A SaaS company cuts its entry-level pricing by 40% to accelerate growth. First order: new signups spike. Second order: existing customers who paid full price discover the new pricing and demand refunds or downgrades — churn accelerates among the segment that was previously most loyal. Third order: the support team, overwhelmed by the refund volume, stops responding to enterprise tickets; two large accounts leave; the sales team, now seeing a damaged reference base, struggles to close at the old price point. The pricing decision looked like a growth move. It was a retention and positioning disaster that only became visible two orders in. The CEO who mapped the second order before deciding would have asked: *how will existing customers react when they see this?* That single question changes the whole shape of the decision.

EXAMPLES

- › *What happens immediately? (First order.) Then what? (Second.) Then what? (Third.)*
- › *How will employees, competitors, regulators, and customers respond to this — beyond the obvious?*
- › *What looks bad now but is good at the second order? What looks good now but bad later?*
- › *Am I stopping at the first consequence because it confirms what I already wanted to do?*



Third-Order Thinking — Map the Full Cascade

Howard Marks (The Most Important Thing, 2011) · Ray Dalio (Principles, 2017) · Systems thinking tradition

Second-order thinking asks *and then what?* once. Third-order thinking asks it twice more — and that second and third iteration is where the most consequential and most invisible consequences live.

Howard Marks: *'First-level thinking is simplistic and superficial. Second-level thinking is deep, complex, and convoluted.'* Third-level thinking adds one more layer: what happens to the broader system — the market, the culture, the competitive landscape, the regulatory environment — once the second-order effects have played out? Most executives reach the second order and stop, believing they have been thorough. The third order is where strategic surprises originate. **A worked example — three orders of**

a talent poaching decision. A CEO recruits a star VP of Engineering from a direct competitor at a 40% salary premium. **First order:** the company gains a high-performer; the competitor loses one.

Second order: the competitor, stung, raises salaries across its engineering team to prevent further losses — and begins actively recruiting from your company in retaliation. Three of your mid-level engineers receive offers within 90 days. **Third order:** compensation benchmarks across the sector shift upward as both companies' moves become known. Smaller competitors can no longer afford

senior engineers at the new market rate. Consolidation accelerates. Two of them are acquired — one by a larger player who now enters your market with engineering depth they didn't have before. The original hire was correct. The third-order consequence reshaped the competitive landscape in a way no one modeled. A CEO who mapped three orders would have asked: *what does this signal to the market, and who benefits from the signal besides us?*

EXAMPLES

- › *First order: what happens immediately? Second order: how does the system react? Third order: how does the landscape shift once those reactions compound?*
- › *Who benefits from the second-order effects of my decision — besides me? If a competitor benefits more than I do at the second order, the decision may be wrong.*
- › *What does this decision signal to the market, to regulators, to talent? Signals are second-order effects that travel faster than the decision itself.*
- › *Where does my chain of consequences become genuinely unpredictable — and have I sized my commitment proportionally to that uncertainty?*
- › *If I mapped this decision's consequences out 18 months, which third-order effect would I most regret not having anticipated?*

Adversarial Thinking — What Does Your Opponent Do Next?

Game theory (von Neumann & Morgenstern, 1944) · Chess strategy · competitive strategy practice

Almost every framework in this guide is single-player: it asks what *you* should do. But most high-stakes decisions happen in a competitive environment where another rational actor will respond to your move. The chess master's reflex — before committing to any move, ask what your opponent's best reply is, then ask if you still want to make the move — is almost entirely absent from executive decision-making. The failure mode: a decision that looks optimal in isolation becomes catastrophic once the competitive response arrives. Amazon enters your market. A key employee joins the company you just passed on acquiring. A regulator responds to the precedent you set. The adversarial thinking discipline is simple: before you decide, steelman your opponent. Give them the best possible response to your move. Then ask whether your decision survives that response. If your plan only works if the other party does nothing — or does exactly what you hope — it is not a plan. It is a wish. The deeper version: map several moves ahead. Your move → their best response → your reply → their reply. At which move does your position become untenable? That is the constraint you need to design around now, not after the game has started.

EXAMPLES

- › *If I make this decision, what is my most capable competitor's best response — not their likely response, their best one?*
- › *Does my plan still work if they make that best response? If not, what needs to change before I commit?*
- › *What does this decision look like from the other side of the table — to the person, company, or regulator who is most threatened by it?*
- › *Am I assuming the other party will be passive, slow, or irrational? What is the cost of being wrong about that?*
- › *Two moves ahead: after their response, what is my reply? Have I thought that far — or am I only thinking one move deep?*

The Great-Good-Ugly Scenario Frame

Warren Berger (The Book of Beautiful Questions, 2018)

When generating options for a decision, project three scenarios: the great outcome, the good (acceptable) outcome, and the ugly (failure) outcome. Berger's point, backed by decision researchers Milkman, Soll, and Payne: *a decision can be no better than the best option under consideration*. Most decision-makers stay binary (yes/no, A/B). The great-good-ugly frame forces at least three distinct option-projections, which exposes which options survive even the ugly scenario — and which only look good when you assume the great one. Works especially well as a group exercise: have each person write their great-good-ugly silently before sharing.

EXAMPLES

- › *For the option I'm considering: what does the **great** outcome look like in 18 months?*
- › *What does the **good** (acceptable but not great) outcome look like?*
- › *What does the **ugly** scenario look like — and is this option still viable in that world?*
- › *Is there an option that's better specifically in the ugly scenario that I haven't considered?*

Asymmetric Bet Sizing — Is the Stake Right?

Warren Buffett (20-slot rule) · Kelly Criterion (J.L. Kelly, 1956) · Nassim Taleb (Antifragile, 2012)

Most decision frameworks ask *whether* to make a bet. Almost none ask whether the size of the bet is right. These are different questions, and confusing them is expensive. A decision that is correct in direction but wrong in magnitude — too much capital, too much time, too many people, too much reputational exposure — can be more damaging than not making it at all. The asymmetric bet principle: **the best bets have limited, defined downside and large, open-ended upside.** Buffett's 20-slot rule: if you imagined you could only make 20 investments in your lifetime, you would not make a small bet on anything uncertain — you would wait for high-conviction, high-upside opportunities and size them properly when they arrived. The Kelly Criterion formalizes this: bet a fraction of your bankroll proportional to your edge divided by your odds. The practical version for executives: before committing, ask not just 'is this right?' but 'if this is right, am I betting enough? And if I'm wrong, is the downside survivable?' The two errors are mirror images: under-betting on a high-conviction opportunity (timidity disguised as prudence) and over-betting on a medium-conviction one (aggression disguised as confidence). Taleb's barbell: keep most exposure in the very safe, take small positions in the very asymmetric — and avoid the middle, where the risk-reward is worst.

EXAMPLES

- › *If this decision is correct, am I betting enough to make it matter — or is my position so hedged it can't generate meaningful upside?*
- › *If this decision is wrong, what is the maximum I lose — and is that survivable? Can I absorb it without it threatening the core?*
- › *What is the ratio of upside to downside here? Is this a 3:1 bet, a 10:1 bet, or a 1:1 bet disguised as something better?*
- › *Am I under-betting because of risk aversion, or because conviction is genuinely low? These require different responses.*
- › *Am I over-betting because the opportunity feels obvious? How many other 'obvious' bets have I seen fail?*
- › *Would I make a larger version of this bet if I had more information — and if so, what's the information I should go get first?*

First Principles Thinking

Aristotle, operationalized by Elon Musk · Weinberg & McCann, Super Thinking, 2019

Most reasoning is by analogy — we see a situation that looks like ones we've seen before and we apply the same pattern. First principles thinking strips that away. You boil the situation down to what you are *actually sure is true*, and reason up from there. Musk's battery example: conventional wisdom said battery packs cost \$600/kWh and always would. First principles: what are the constituent materials? What do they cost on the London Metal Exchange? \$80/kWh. The conventional wisdom was just a number people had been repeating, not a law of physics. For CEO decisions: when the room agrees that 'this is just how it works in our industry,' that's usually the signal to ask what's actually structurally true versus what's inherited assumption.

EXAMPLES

- › *What am I treating as a fixed constraint that might actually be a convention, a habit, or received wisdom?*
- › *If I stripped away what 'everyone knows,' what would actually remain as demonstrably true?*
- › *What are the constituent parts of this problem — the raw materials — rather than the assembled product I've been handed?*
- › *If someone with no industry experience but unlimited intelligence looked at this, what would they conclude?*

Known Unknowns / Unknown Unknowns

Luft & Ingham (1955) · Rumsfeld (2002) · Weinberg & McCann, Super Thinking, 2019

Most decision processes only test the **known unknowns** — risks you're aware of but haven't resolved. The genuinely dangerous category is the **unknown unknowns**: things you don't know that you don't know. You can't Google for them. You can only create conditions where they're more likely to surface: diverse perspectives, pre-mortems, crystal ball exercises, scenario analysis. The pre-mortem in Stage 06 is specifically designed to convert unknown unknowns into known unknowns before you commit. The failure mode is believing your risk register is complete when it's actually just a list of the risks that were easy to name.

EXAMPLES

- › *What's on my risk list that I know about and haven't resolved yet? (Known unknowns — at least name them.)*
- › *What risks are not on my list — not because they don't exist but because I haven't thought to look for them?*
- › *Who might know something I don't know I need to know — and have I asked them?*
- › *If the unknown unknown that sinks this decision surfaced in 12 months, what domain would it come from?*

Commitment Escalation — When Does a Pilot Become a Trap?

Barry Staw (escalation of commitment research, 1976) · Sunk Cost Bias · Roger Miller & Donald Lessard

The sunk cost bias describes what happens after a commitment goes wrong: you keep investing because of what's already been spent. Commitment escalation is the earlier, subtler trap: a series of individually reasonable decisions that, in aggregate, converts a Type 2 reversible choice into a Type 1 irreversible one — *before you notice the transition has happened*. The mechanism: each small step seems justified. A pilot launches. Early results are mixed but not bad. The team doubles down. A vendor is signed. A hire is made. Press coverage comes out. A board update is given. At each step, the cost of reversing increases — financially, reputationally, and psychologically. By the time the data clearly says this isn't working, reversal feels catastrophic when it is merely expensive. Barry Staw's research identified the key driver: decision-makers who feel personally responsible for the initial commitment escalate far more than those who inherited the decision. The question is not whether you should persist through difficulty — you should. The question is whether you are persisting based on new evidence or defending a position because it's yours. The practical diagnostic: if you would not make this commitment fresh today, with full knowledge of what you now know, the case for continuing rests on sunk cost, not on merit. That is the line.

EXAMPLES

- › *If I were arriving at this decision today — no prior investment, no prior announcement — would I still make it?*
- › *At what point in this project did it change from reversible to difficult-to-reverse? Did I notice that transition at the time?*
- › *Am I continuing because the evidence supports it, or because I am the one who started it?*
- › *What would it actually cost to stop now — in money, relationships, and reputation — versus what it will cost to continue if this doesn't work?*
- › *Have I set a specific kill criterion in advance, or am I evaluating this decision on a rolling basis with no pre-committed exit point?*
- › *Is the team's momentum on this decision an asset — or is it the thing that's preventing me from seeing clearly?*

Goodhart's Law

Charles Goodhart (1975) · Marilyn Strathern (1997) · via Weinberg & McCann, *Super Thinking*, 2019

When a measure becomes a target, it ceases to be a good measure. The most important implication for executive decisions: any metric you set as a tripwire or success criterion will eventually be optimized *for* in ways that decouple it from what you actually care about. Revenue targets get hit by pulling forward deals. Uptime metrics get gamed by excluding certain incident types. NPS scores get inflated by selectively surveying happy customers. This doesn't mean don't set metrics — it means build in periodic audits of whether the metric still measures what it was supposed to measure, and expect that it will drift.

EXAMPLES

- › *What am I measuring as a proxy for what I actually care about — and could that proxy be gamed?*
- › *If someone on my team wanted to hit this target without actually achieving the underlying goal, how would they do it?*
- › *How will I know if the metric has drifted from the thing it was supposed to represent?*
- › *Is there a second metric that would make gaming the first one visible?*

The Cobra Effect

Named after a colonial-era British policy in India · via Weinberg & McCann, Super Thinking, 2019

When an attempted solution makes the problem worse. The British colonial government in India, alarmed by cobra snakes, offered a bounty for every dead cobra. Enterprising citizens bred cobras to collect the bounty. When the program was cancelled, the breeders released their worthless cobras, dramatically increasing the cobra population. The cobra effect is Goodhart's Law plus perverse incentives plus Chesterton's Fence combined. For a CEO: before implementing any solution, ask what the most rational response of a self-interested person would be to the incentives you're creating — and whether that response advances or undermines your actual goal.

EXAMPLES

- › *Who benefits from the incentive structure I'm creating — and would their self-interested response actually advance my goal?*
- › *What would someone do to exploit or game this policy — and is that worse than the problem I'm solving?*
- › *Am I solving the symptom in a way that makes the underlying problem harder to see or address?*
- › *Does fixing this decision create a downstream problem I haven't accounted for?*

Hanlon's Razor

Named after Robert Hanlon · via Weinberg & McCann, Super Thinking, 2019

Never attribute to malice that which is adequately explained by carelessness. For a CEO, this applies most acutely to information hygiene: when a key fact is missing from a briefing, when a team member delivers a misleading report, when someone's framing seems designed to lead you somewhere — the first question is usually not 'are they lying?' but 'did they not realize this mattered?'. Hanlon's Razor is also a leadership disposition: most people are not trying to deceive you. They're busy, under pressure, and shaped by their own frame. Assuming malice burns trust; assuming carelessness leaves room for correction. Apply it to yourself too: before concluding you were misled, ask whether the explanation was just incomplete — then ask why.

EXAMPLES

- › *Before assuming this was intentional, what's the simplest explanation that doesn't require bad faith?*
- › *If this person made an error rather than a choice, what would I do differently?*
- › *Am I reacting to what happened or to the story I've built about why it happened?*
- › *What question would I ask if I genuinely assumed good intent — and should I start there?*

Black Swan / Fat Tails

Nassim Taleb (The Black Swan, 2007) · Weinberg & McCann, Super Thinking, 2019

A black swan is an extreme, high-impact event that was considered nearly impossible — until it happened. The term comes from the European belief, held for centuries, that black swans didn't exist. The operative insight for decision-makers: many situations that look like normal-distribution risks are actually *fat-tailed* — the extreme events occur more often than probability models suggest. This is especially true for complex, interconnected systems (financial markets, supply chains, geopolitics). For a CEO committing to a major irreversible decision: the scenario you've dismissed as 'essentially impossible' may have a 5–10% real probability. The question is not whether black swans happen — they do — but whether your decision survives one.

EXAMPLES

- › *What's the scenario I've dismissed as essentially impossible — and what's its actual probability if I force myself to estimate?*
- › *If the most extreme plausible bad outcome occurred, would this decision still be recoverable?*
- › *Am I treating a complex, interconnected system as if it follows predictable, normal-distribution risk?*
- › *Have I stress-tested this decision against a fat-tail scenario — not just the expected one?*

OODA Loop — Orient Before You Decide

Col. John Boyd, U.S. Air Force (1960s–1990s) · via Shane Parrish, Farnam Street, 2018

Boyd's OODA loop — Observe, Orient, Decide, Act — was developed for fighter pilots making life-or-death decisions in seconds. The full loop is well known. What is consistently underestimated is the second step: **Orient**. Boyd called it the *schwerpunkt* — the main emphasis — and it is the step most executives skip entirely. To orient is to ask: *which mental model am I using to interpret what I just observed — and is it the right one for this situation?* Boyd identified four barriers that corrupt orientation: cultural traditions (inherited assumptions), genetic heritage (cognitive hard-wiring), analytical habits (the frameworks you always reach for), and the pace of new information (which makes it hard to update the model in time). His remedy: 'destructive deduction' — actively dismantling the wrong model before building a better one. The executive failure mode this addresses is not missing information. It is applying the wrong interpretive frame to correct information and reaching the wrong conclusion with high confidence. Boyd's second insight — equally important — is the value of **tempo**: cycling through the loop faster than the situation changes gives you initiative. For a CEO, this means the quality of your decisions compounds when you update quickly, explicitly, and without ego. The pilot who can re-orient in two seconds beats the one who needs ten — even if the second pilot's eventual analysis is more thorough.

EXAMPLES

- › *Before deciding: which mental model am I using to interpret this situation — and did I choose it deliberately, or did it arrive by default?*
- › *What would someone with a completely different professional background see in this same situation that I'm not seeing?*
- › *Is the frame I'm using the one that served me well last time — and is this situation actually similar enough for that to be valid?*
- › *Am I cycling through Observe → Orient → Decide → Act fast enough to stay ahead of the situation — or is the situation updating faster than I am?*
- › *What assumption would I have to dismantle — destroy, not just question — for a different decision to become obvious?*

Inversion — Think Backward

Charlie Munger (Poor Charlie's Almanack, 2005) · Carl Jacobi, 19th c. mathematician · Jasmine Bina, Concept Bureau, 2025

Instead of asking how to achieve the outcome, ask how to guarantee failure. Then avoid doing those things. Munger's formulation: 'Invert, always invert.' The mathematician Carl Jacobi solved hard problems by reversing them — 'man muss immer umkehren' (always invert). The technique works because forward thinking is optimistic and selective; backward thinking is brutally honest. When you ask 'how do we succeed?' your mind reaches for hopeful plans. When you ask 'how would we guarantee we fail?' your mind immediately surfaces the real constraints, the actual risks, and the things everyone privately knows but doesn't say out loud. Inversion complements the pre-mortem (Stage 06) but operates differently: the pre-mortem imagines a future failure and works backward to causes; inversion starts with the failure modes and engineers their prevention into the plan before committing. Applied to decisions: before you decide how to succeed, fully specify how to fail. The gap between that list and your current plan is the risk you haven't designed around. The Dallas–Fort Worth Airport example: passengers complained about baggage wait times. Inversion — 'how would we guarantee the wait feels intolerable?' — pointed directly to the solution: make baggage claim farther from gates so passengers walk more and wait less. Complaints disappeared without any change to baggage delivery speed.

EXAMPLES

- › *How would I guarantee this decision fails? Be specific — name the mechanisms, not the feelings.*
- › *Which of those failure modes is already present in the current plan — and have I designed around it?*
- › *What would the worst version of this decision look like in 18 months? Am I doing anything that points in that direction?*
- › *If I were trying to sabotage this decision from the inside, what would I do — and is anyone already doing it?*
- › *What does the inverted question reveal that the forward question missed?*

Local vs. Global Maxima

Mathematics / optimization theory · Jasmine Bina, Concept Bureau, The Strategist's Playbook, 2025

A local maximum is the best solution within a limited search space. A global maximum is the best solution across the entire possibility space. The problem: from inside a local maximum, it is impossible to see the global one — and the path to the global maximum always runs through a valley, a period of worse performance before better. This is the mechanism behind the innovator's dilemma: companies optimize their current position so efficiently that they cannot see — or cannot stomach — the descent required to reach the next peak. For CEO decisions, the local maximum trap shows up in two forms. First: **the optimization trap** — making a good decision for the current game when the game itself is changing. A great retail strategy in 2008. A great on-premise software roadmap in 2012. Right answer, wrong hill. Second: **the adjacent hill problem** — the next opportunity is visible, and better, but requires temporarily giving up performance on the current one. Leaders who cannot tolerate the valley cannot reach the next peak. The diagnostic question: are you optimizing for today's rules, or positioning for tomorrow's? These are different decisions. The first feels safer. The second is usually more correct.

EXAMPLES

- › *Am I optimizing for the current rules of the game — or for the rules that will govern the game in three years?*
- › *Is this decision the best move on the current hill, or the best move for reaching the next hill — and are those the same thing?*
- › *What would I have to give up, or tolerate being worse at, to reach a meaningfully better position? Am I willing to go through the valley?*
- › *Is our current strength — the thing we're optimizing — also the thing that makes it hardest to see the next opportunity?*
- › *What does the global maximum look like from the outside — and why does the path there feel like the wrong direction from the inside?*

APPENDIX F

Key insights from the source library

The ideas in Before You Commit were distilled from a specific set of books. Below are the most important insights from each — not summaries, but the sharpest, most non-obvious ideas worth carrying into any major decision. Use this appendix as a pre-decision checklist: scan it before committing.

Decisive: How to Make Better Choices in Life and Work

Chip & Dan Heath, 2013 — Source of Stage 03 (Widen) and Stage 05 (Attain Distance)

- › **The spotlight effect.** We give too much weight to information in front of us and too little to what's just offstage. Every decision has an implicit spotlight — and what's outside it is usually where the real answer is.
- › **Ooching before committing.** Rather than analyzing your way to certainty, run a small experiment that generates real information. If you think the wireless-sensor market is promising, test it first. Don't predict — prototype. Most decisions that feel un-oochable actually aren't.
- › **The vanishing-option test.** If your preferred option suddenly became unavailable, what would you do instead? That answer reveals whether you've genuinely considered alternatives — or just looked for justifications.
- › **Inside view vs. outside view.** The inside view draws from the unique features of your situation; the outside view ignores them and looks at base rates. You are 100% more confident in your inside-view prediction than the data warrants. Force the outside view every time.
- › **Honor your core priorities.** Most urgent decisions crowd out the important ones. The Heath brothers found that many executives spent zero time on their stated core priorities across an entire week — not because they chose not to, but because urgencies filled every gap. Core priorities need to be enshrined, not just declared.
- › **Bookend the future.** Don't just plan for one future. Bracket the range: what's the realistic best case? The realistic worst case? The decision that's good in both scenarios is usually better than the one optimized for the expected case.
- › **Set a tripwire.** A pre-committed, specific signal that forces a re-evaluation. Set in cold blood. Not a vague intention to 'check in' — a specific number, a specific date, a specific event. Van Halen's brown M&M clause is the canonical example.

How to Decide

Annie Duke, 2020 — Source of Appendix D (Decision Journal) and Stage 09 (Calibration)

- › **Resulting.** The most pervasive cognitive trap in executive decision-making: judging decision quality by outcome quality. You run a red light and get through safely — great decision? No. Resulting corrupts the feedback loop. The antidote is explicitly separating process quality from outcome quality — every time, including when evaluating others.
- › **The Happiness Test.** Before investing time in a decision, ask: in a year, will I care how this turned out? If not — if the decision passes the Happiness Test, meaning it's genuinely low-stakes — go fast. The Happiness Test exposes how much CEO time goes to decisions that really don't matter at their altitude.
- › **Inoculation.** Before committing, write down how you'll evaluate this decision independent of the outcome. This is the only way to protect against resulting after the fact. The decision journal is the physical form of inoculation.
- › **Resulting is contagious.** When you evaluate your team's decisions by outcomes, you teach them to make safer (less value-generating) decisions. They learn to avoid decisions that could go wrong visibly, not to make better decisions. The culture of a good decision-making organization rewards good process even when the outcome is bad.
- › **Backcasting.** Imagining success and working backward tells you what to build in — not just what to protect against. Pre-mortem (failure) + backcast (success) together cover the full range. Most leaders do neither systematically.

Solvable: A Simple Framework for Complex Problems

Arnaud Chevallier & Albrecht Enders, 2022 — Source of MECE and LEAD techniques in Appendix E

- › **The FrED framework.** Frame → Explore → Decide. The most disciplined problem-solving structure in the literature. Frame first (the quest), explore second (the alternatives), decide third. Most executives jump to Decide before completing Frame — which is why the wrong problem gets solved with great rigor.
- › **The Hero-Dragon-Quest frame.** A good problem frame has four components: who is deciding (hero), what they want to achieve (treasure), what stands in the way (dragon), and the overarching question (quest). If you can't articulate all four, you haven't finished framing. The dragon is the most commonly skipped — and the most important.
- › **MECE: the option-space discipline.** Mutually Exclusive means your options don't overlap. Collectively Exhaustive means they cover the full solution space. Most executive option-sets fail CE — they have large blind-spot gaps. The gap is usually where the best answer lives.
- › **The why-map.** Before identifying root causes, map all possible causes visually — don't start with a single root-cause hypothesis. The why-map is a tree of plausible causes; each branch is a hypothesis to test, not a conclusion. Most root-cause analyses tunnel prematurely on the first plausible explanation.
- › **LEAD your hypotheses.** Locate relevant evidence before looking (to avoid cherry-picking). Evaluate quality of each source. Synthesize the body of evidence. Decide. The critical discipline: favor opposing evidence. Supporting evidence builds confidence; opposing evidence builds accuracy.
- › **Weighted criteria.** Once you have options, evaluate them against explicit, weighted criteria — not just intuition. Force yourself to state what matters and how much before you evaluate. This breaks anchoring: without pre-stated weights, you subconsciously adjust importance to match your preferred option.

The Book of Beautiful Questions

Warren Berger, 2018 — *Source of leadership questioning questions in Appendix E*

- › **'And what else?' — the AWE question.** Executive coach Michael Bungay Stanier calls this the best coaching question in the world. After someone answers a question, ask: 'And what else?' three times. It pushes beyond top-of-mind answers, keeps you in listening mode, and suppresses 'the advice monster' — the impulse to move to solutions before the problem is fully surfaced.
- › **WAIT: Why Am I Talking?** Psychologist Ronald Siegel's self-check for executives in meetings. Before speaking, ask yourself WAIT. Most executive input in a discussion is premature opinion that shuts down the room. The leader who talks last gets better information from the room than the one who talks first.
- › **The confidence-to-be-humble paradox.** Berger's research on the best leaders consistently found that they had the confidence to ask the questions others were afraid to ask — and the humility to accept that the answers might change their view. These two traits are usually presented as opposites; in practice they compound each other.
- › **Questioning down is harder than questioning up.** Getting genuine answers from people who report to you is structurally harder than getting answers from peers. People self-censor upward. The most useful question a CEO can ask a direct report is not 'what do you think?' but 'what would someone who disagreed with this say?'
- › **Five enemies of questioning:** Fear (of looking naïve), knowledge (expertise that makes you stop asking), bias (assumptions that shape what you notice), hubris (certainty that closes inquiry), and time (the pressure that removes reflection). Of these, hubris is the one most correlated with executive altitude — and most resistant to self-correction.

Farnam Street / Shane Parrish

Shane Parrish (fs.blog) · *mental models, second-order thinking, inversion* · Nir Eyal (nirandfar.com) on *present-bias*

- › **Mental models reduce blind spots.** The models in your head determine which information you notice and which you ignore. The CEO who only has a sales model will miss the churn signal. Adding mental models from different disciplines — systems thinking, thermodynamics, reciprocity, inversion — is the most durable investment in decision quality available.
- › **Second-order thinking: 'And then what?'** Most decisions look good at first order. The second order is where competitive response, regulatory reaction, cultural backlash, and system-level effects live. The third order is where compounding happens — in both directions. Parrish: failing to consider second-order consequences is especially deadly when the first order confirms your biases.
- › **Inversion: think backward.** Instead of asking how to achieve the outcome, ask how to guarantee it fails. Remove the failure modes. What you're left with is a more robust path to success than direct forward planning produces. Munger: 'Invert, always invert.'
- › **Socratic probing has five categories.** Each targets a different cognitive blind spot: Clarification (what do you mean exactly?), Probing Assumptions (what are you assuming?), Probing Evidence (is there reason to doubt this?), Questioning Perspectives (what would a disagreeer say?), Probing Implications (then what would happen?). Applied systematically to a team's recommendation, these five categories expose more flaws than any single question.
- › **Hyperbolic discounting — the present-bias mechanism.** Nir Eyal: we systematically overvalue immediate rewards relative to future ones — even when we know better. For executive decisions, this means the option that feels best today is usually the one most infected by present-bias. The 10-10-10 technique (how will you feel about this in 10 minutes, 10 months, 10 years?) is the fastest antidote.
- › **Avoiding stupidity is easier than seeking brilliance.** Parrish's restatement of Munger: if you're focused on avoiding the obvious failure modes — resulting, narrow framing, confirmation, sunk cost — you will make better decisions than most of your peers without doing anything else. World-class decision-making is mostly about not being stupid in the predictable ways.

Money and Love: An Intelligent Guide to the World's Two Biggest Decisions

Myra Strober & Abby Davisson · Crown Currency, 2023 — Stanford GSB 'Work and Family' course distilled

- › **The premise most decision guides ignore.** Career decisions and relationship decisions are not separate. Every major professional choice — where to live, whether to take a promotion, how much to travel — is simultaneously a relationship decision, and vice versa. Treating them in isolation produces locally optimized, globally incoherent choices. When the decision on your desk looks purely professional, ask who else's life it reconfigures.
- › **Mimetic desire corrupts framing.** René Girard's term: we desire what the people around us desire, often without realizing it. For executives, this means the options on the table — and the criteria for evaluating them — are often inherited from peers, industry norms, or institutional pressure rather than chosen deliberately. Before clarifying what you want, ask: *whose wanting am I enacting?* This is a version of Stage 08's identity interrogation, applied upstream of the decision.
- › **The decision that keeps not getting made.** The most consequential decisions — about partners, structure, location, organizational shape — are the ones deferred most systematically. The cost of not deciding accrues silently. Not-deciding is a decision with compounding consequences. A forcing function has to be created deliberately; it will not arrive on its own.
- › **Consequences across three time horizons.** For any major decision, map consequences explicitly at: the next few months (immediate), six months to two years (short term), and beyond two years (long term). What looks costly in the immediate often becomes neutral or positive in the long term. What looks neutral in the immediate is often where the compounding damage or gain lives. Stage 06 of this guide does this for failure modes; extend the habit to all consequence types.
- › **The whole-life frame.** Strober and Davisson's deepest instruction: the question is never only what this decision does for your career, or only what it does for your relationship. The question is what it does for the shape of your life over ten years. Most executive decision processes have no mechanism for asking this. The absence of the mechanism is itself a decision.

Most readers will benefit more from rereading two of these books carefully than from reading all of them once.

*The best decision-makers are not the ones who are right most often.
They are the ones who notice they are wrong soonest,
and adjust without ego.*

APPENDIX G

The reading list

Curated, not exhaustive. The decision-making literature is enormous and most of it is restatement of a small number of original ideas. The books and essays below are the ones that add something the others don't — each grouped by what you'd reach for it for, with a sharp one-line take on why it matters. ★ marks the four that, if you read nothing else, are sufficient.

On how decisions actually go wrong

The cognitive science underneath every other book on this list

★ Thinking, Fast and Slow

Daniel Kahneman, 2011

The foundation. System 1 and System 2, the major biases, what 40 years of behavioral economics actually established. Long, but everything else assumes you've read it. Read the chapters on overconfidence, planning fallacy, and loss aversion first — they're the ones that bite executives most often.

· Sources of Power: How People Make Decisions

Gary Klein, 1998 (MIT Press, 2017 ed.)

The counterpoint to Kahneman — and the book that originated the pre-mortem. Klein studied firefighters, military commanders, and nurses making real decisions under time pressure and found they *don't* compare options: experts recognize situations as familiar and mentally simulate the first viable course of action. Read it to understand where your fast intuition is reliable versus where it needs deliberate override.

· Predictably Irrational

Dan Ariely, 2008

Faster and more entertaining than Kahneman, with experiments that stick. Less rigorous, but better for getting the patterns into your bones.

· Noise: A Flaw in Human Judgment

Kahneman, Sibony, Sunstein, 2021

Kahneman's follow-up. The argument: bias gets all the attention, but *noise* — random variation in judgment from the same person across occasions — is just as costly and almost never measured.

On making decisions better

Practical frameworks you can actually apply tomorrow

★ Decisive: How to Make Better Choices in Life and Work

Chip & Dan Heath, 2013

The WRAP framework — Widen options, Reality-test, Attain distance, Prepare to be wrong — is the closest thing the field has to a complete operational guide. The 'ooching' concept (test before committing), the inside-view vs. outside-view distinction, and the vanishing-option test are the three most immediately applicable ideas in it. If you only read one decision book, read this.

★ Thinking in Bets

Annie Duke, 2018

Distinguishes decision quality from outcome quality, introduces the decision journal, and explains why poker players are systematically better calibrated than executives. Essential for anyone who confuses being right with getting lucky.

· How to Decide

Annie Duke, 2020

Duke's companion workbook — more structured and practical than *Thinking in Bets*. Introduces 'resulting' (the trap of judging past decisions by their outcomes) and 'inoculation' (pre-committing to how you'll evaluate a decision before the outcome arrives). The Happiness Test for sorting low-impact from high-impact decisions is alone worth the read.

- **Solvable: A Simple Framework for Complex Problems**

Arnaud Chevallier & Albrecht Enders, 2022

The FrED framework (Frame, Explore, Decide), MECE discipline, the why-map, LEAD evidence testing, and weighted criteria. The most disciplined structured problem-solving process available in book form. Pairs well with Decisive.

- **How Will You Measure Your Life?**

Clayton Christensen, 2012

Less about decisions in the moment, more about the cumulative shape of decisions over decades. The chapter on marginal thinking — "just this once" — explains more career-ending CEO errors than any other single idea.

- **Superforecasting**

Philip Tetlock & Dan Gardner, 2015

The book on calibration. If you want to actually get better at probability estimates over time, this is the protocol. Pair with the decision journal.

On asking better questions

The skill upstream of every good decision

★ A More Beautiful Question

Warren Berger, 2014 (10th anniv. ed. 2024)

The source of Appendices A and E in this guide. The argument: most breakthroughs trace to a single well-formed question, and most leaders have forgotten how to ask one. Read it twice — once for the patterns, once for the stories you'll want to retell.

- **The Book of Beautiful Questions**

Warren Berger, 2018

Berger's companion volume organized as a working reference: 400+ questions for decisions, creativity, leadership, and connection. The leadership section alone contains more operational prompts than most executive coaching engagements.

- **What's Your Problem?**

Thomas Wedell-Wedellsborg · Harvard Business Review Press, 2020

Source of the elevator-with-mirrors story in Stage 01, the Bright Spots technique, and the Five Reframing Moves in Appendix E. Read it before any major framing exercise — shorter and more practical than most books on the topic.

- **Questions Are the Answer**

Hal Gregersen, 2018

MIT's research on 'question bursts' — generating 20+ questions in 4 minutes to break a stuck decision. The empirical case behind the technique.

On deciding from the top

What changes when the decision is yours alone

★ High Output Management

Andy Grove, 1983

Still the best book on operational executive judgment, forty years later. Grove's framework for decisions made by groups — and his test for whether your meetings are decision-making or decision-faking — is more useful than anything written since.

- **The Hard Thing About Hard Things**

Ben Horowitz, 2014

What no business school teaches: how to fire your friend, demote your co-founder, decide in the absence of good options. Voice is sharper than most CEO books. Read it for the section on "the struggle."

- **Only the Paranoid Survive**

Andy Grove, 1996

Grove's book on strategic-inflection-point decisions — the moments when the rules of your industry are changing and you have to decide whether to bet the company on the new game.

- **What You Do Is Who You Are**

Ben Horowitz, 2019

On the decisions that build culture, which compound more than any single operational call. Underrated chapter on Toussaint Louverture as a decision-maker.

On strategy and base rates

Resisting the inside view of your own situation

- **Good Strategy / Bad Strategy**

Richard Rumelt, 2011

The clearest writing on what strategy actually is, and why most documents labeled "strategy" aren't. Chapter on "the kernel" is required reading before any strategic Type 1 decision.

- **Super Thinking: The Big Book of Mental Models**

Gabriel Weinberg & Lauren McCann, 2019

300+ mental models from every discipline. Source of Goodhart's Law, the Cobra Effect, Hanlon's Razor, Known/Unknown Unknowns, First Principles, and Black Swan / Fat Tails in Appendix E. Not a cover-to-cover read — a reference to consult when a decision feels stuck and you can't name why.

- **Seeking Wisdom: From Darwin to Munger**

Peter Bevelin, 2007

Charlie Munger's mental models, organized. Dense, not narrative — closer to a reference manual than a book. Most useful as a lookup when you suspect a specific bias is operating but can't name it.

- **Poor Charlie's Almanack**

Charles Munger, 2005

The source material for Bevelin's distillation. Munger's lectures on decision-making, inversion, and the psychology of human misjudgment. Worth owning physically — you'll return to it.

Essays worth bookmarking

What books haven't covered, or haven't covered well

- **Performing a Project Premortem**

Gary Klein, Harvard Business Review, 2007

The original five-page article that introduced the pre-mortem to business. The single highest-leverage decision technique in this guide, by the person who developed it.

- **Before You Make That Big Decision...**

Kahneman, Lovallo, Sibony, HBR, 2011

A 12-question checklist for evaluating someone else's recommendation — exactly the situation most of this guide addresses from the other side.

- **The Cognitive Style of PowerPoint**

Edward Tufte, 2003

How slide decks systematically degrade the decisions they're meant to support. Read it the next time you're asked to decide based on a deck.

- **Type 1 and Type 2 Decisions**

Jeff Bezos, 1997 Amazon shareholder letter

The one-page formulation that anchors Stage 2 of this guide.

- **The Scout Mindset**

Julia Galef, 2021

Calibration and intellectual honesty. The distinction between soldier mindset (defending a position) and scout mindset (mapping the terrain) is operational.

Practitioner writing worth reading

People who have done it, writing about it — shorter, faster, no academic hedging

- **Being Strategic: How to Analyze Problems**

Joseph Gefroh, substack.com/p/being-strategic-how-to-analyze-problems, 2025

Source of Stage 01's multi-frame diagnosis and Chesterton's Fence technique. The CRO churn story alone is worth the read. Part of his *Being Strategic* series for senior managers moving into executive roles.

- **How to Make Smart Decisions Without Getting Lucky**

Shane Parrish · Farnam Street (fs.blog)

Second-order thinking, inversion, avoiding stupidity as the core discipline. One of the best free resources in the decision-making space.

- **James Clear's Decision-Making Articles**

James Clear · jamesclear.com

The decision-making and mental-models articles specifically (not the habits content). Buffett's 20-slot rule, inversion, motion vs. action. Good for sharing with teams who won't read the books.

On decisions that span work and life

A different angle — what the standard decision literature leaves out

- **Money and Love: An Intelligent Guide to the World's Two Biggest Decisions**

Myra Strober & Abby Davisson · Crown Currency, 2023

Strober spent forty years teaching Stanford GSB's Work and Family course. Her central claim: professional decisions and personal decisions are not separable, and treating them as if they are produces incoherent choices. The 5Cs framework (Clarify, Communicate, Consider, Check in, Consequences) is a clean decision process for any high-stakes life decision. Read it for the whole-life frame — the reminder that the question is never only what a decision does for your career, but what it does for the shape of your life over ten years.

Most readers will benefit more from rereading two of these books carefully than from reading all of them once.

*The best decision-makers are not the ones who are right most often.
They are the ones who notice they are wrong soonest,
and adjust without ego.*

APPENDIX H

All sources

Every book, article, framework, and practitioner source drawn on in *Before You Commit*, listed by category. Where a source shaped a specific stage, technique, or appendix, that is noted.

Books — Core Decision-Making

› Decisive: How to Make Better Choices in Life and Work

Chip & Dan Heath · Harvard Business Review Press, 2013

Source of WRAP framework, ooching, inside/outside view, vanishing-option test, tripwires. Stage 03, 05, Appendix E.

› Thinking in Bets

Annie Duke · Portfolio/Penguin, 2018

Source of resulting, decision journal, calibration. Appendix B, D, Stage 09.

› How to Decide

Annie Duke · Portfolio/Penguin, 2020

Source of Happiness Test, inoculation, backcasting. Appendix E, G.

› Sources of Power: How People Make Decisions

Gary Klein · MIT Press, 1998 (2017 ed.)

Inventor of the pre-mortem and crystal ball technique. Stage 06. Appendix F.

› Solvable: A Simple Framework for Complex Problems

Arnaud Chevallier & Albrecht Enders · FT Press, 2022

Source of FrED framework, Hero-Dragon-Quest (Appendix E card), MECE, why-map, LEAD, weighted criteria. Appendix E, F, G.

› Thinking, Fast and Slow

Daniel Kahneman · Farrar, Straus and Giroux, 2011

Foundation for bias appendix. Appendix A.

› Noise: A Flaw in Human Judgment

Kahneman, Sibony, Sunstein · Little, Brown Spark, 2021

Appendix F.

› Predictably Irrational

Dan Ariely · Harper, 2008

Appendix F.

› Superforecasting

Philip Tetlock & Dan Gardner · Crown, 2015

Calibration and probability estimation. Appendix D, F.

Books — Questions & Thinking

› A More Beautiful Question

Warren Berger · Bloomsbury, 2014 (10th anniv. ed. 2024)

Source of Why→What If→How, constraint-removal question, Polaroid/Airbnb/Netflix stories. Appendix A, E, G.

› The Book of Beautiful Questions

Warren Berger · Bloomsbury, 2018

Source of AWE question, WAIT, five enemies of questioning, Great-Good-Ugly frame. Appendix E, G.

› **Questions Are the Answer**

Hal Gregersen · HarperBusiness, 2018

Question bursts technique. Appendix F.

› **What's Your Problem?**

Thomas Wedell-Wedellsborg · Harvard Business Review Press, 2020

Source of elevator-with-mirrors story, Bright Spots technique, Five Reframing Moves, baked-in-solution prompt. Stage 01, Appendix E, F.

› **Super Thinking: The Big Book of Mental Models**

Gabriel Weinberg & Lauren McCann · Portfolio/Penguin, 2019

Source of Goodhart's Law, Cobra Effect, Hanlon's Razor, Known/Unknown Unknowns, First Principles, Black Swan/Fat Tails. Appendix E.

Books — Executive Leadership

› **High Output Management**

Andy Grove · Random House, 1983

Operational executive judgment. Appendix F.

› **The Hard Thing About Hard Things**

Ben Horowitz · Harper Business, 2014

Appendix F.

› **Only the Paranoid Survive**

Andy Grove · Currency, 1996

Strategic inflection points. Appendix F.

› **What You Do Is Who You Are**

Ben Horowitz · Harper Business, 2019

Appendix F.

Books — Strategy & Mental Models

› **Good Strategy / Bad Strategy**

Richard Rumelt · Crown Business, 2011

Appendix F.

› **Seeking Wisdom: From Darwin to Munger**

Peter Bevelin · PCA Publications, 2007

Mental model reference. Appendix F.

› **Poor Charlie's Almanack**

Charles Munger · Donning, 2005

Inversion, mental models, psychology of misjudgment. Appendix F.

› **The Scout Mindset**

Julia Galef · Portfolio/Penguin, 2021

Calibration, soldier vs. scout mindset. Appendix F.

› **How Will You Measure Your Life?**

Clayton Christensen · Harper Business, 2012

Marginal thinking, cumulative decision patterns. Appendix F.

› **The Little Book of Big Decision Models**

James McGrath · FT Press, 2016

Source of RAPID framework (Rogers & Blenko). Appendix E.

› **Money and Love: An Intelligent Guide to the World's Two Biggest Decisions**

Myra Strober & Abby Davisson · Crown Currency, 2023

Source of 5Cs framework, whole-life decision frame, mimetic desire in option-generation, three-horizon consequence mapping. Appendix F, G.

Practitioner Articles & Essays

› **Performing a Project Premortem**

Gary Klein · Harvard Business Review, 2007

Original five-page pre-mortem article. Stage 06, Appendix F.

› **Before You Make That Big Decision...**

Kahneman, Lovallo & Sibony · Harvard Business Review, 2011

12-question bias checklist for evaluating recommendations. Appendix F.

› **Being Strategic: How to Analyze Problems**

Joseph Gefroh · substack.com/p/being-strategic-how-to-analyze-problems, 2025

Source of multi-frame diagnosis, Chesterton's Fence application, CRO churn reframe story. Stage 01, Appendix E, G.

› **Type 1 and Type 2 Decisions**

Jeff Bezos · 1997 Amazon shareholder letter

Source of one-way/two-way door distinction. Stage 02, Appendix F.

› **The Cognitive Style of PowerPoint**

Edward Tufte · Graphics Press, 2003

Appendix F.

› **How to Make Smart Decisions Without Getting Lucky**

Shane Parrish · Farnam Street (fs.blog), ongoing

Source of second-order thinking, inversion, avoiding stupidity framing. Appendix E, G.

› **Decision Making: How to Make Better Life Choices**

Nir Eyal · nirandfar.com, ongoing

Source of hyperbolic discounting mechanism. Appendix G.

Frameworks & Historical Sources

› **Socratic Questioning Taxonomy**

Socrates (469–399 BC), operationalized in clinical and educational practice

Five-category framework: Clarification, Probing Assumptions, Probing Evidence, Questioning Perspectives, Implications. Appendix E.

› **De Bono's Six Thinking Hats**

Edward de Bono · Little, Brown, 1985

Parallel thinking framework. Appendix E.

› **Chesterton's Fence**

G.K. Chesterton · 'The Thing', 1929

Don't remove a constraint until you understand why it exists. Stage 01, Appendix E.

› **Goodhart's Law**

Charles Goodhart, 1975; formalized by Marilyn Strathern, 1997

When a measure becomes a target it ceases to be a good measure. Appendix E.

› **The Cobra Effect**

Named after British colonial India policy; popularized by Horst Siebert, 2001

When a solution makes the problem worse. Appendix E.

› **Hanlon's Razor**

Named after Robert Hanlon; attributed in various forms to Goethe and Napoleon

Never attribute to malice what is adequately explained by carelessness. Appendix E.

› **The Pre-Mortem / Crystal Ball**

Gary Klein & colleagues, Klein Associates, 1985–1990s

Stage 06, Appendix E.

› **RAPID Decision Framework**

Paul Rogers & Marcia Blenko · Bain & Company

Decision role clarity: Recommend, Agree, Perform, Input, Decide. Appendix E.

› **MECE: Mutually Exclusive, Collectively Exhaustive**

McKinsey & Company (attributed to Barbara Minto, 1970s)

Appendix E.

› **How Might We...?**

Min Basadur, Procter & Gamble, 1970s; adopted by IDEO (Tim Brown)

Appendix E.

› **Bezos Type 1 / Type 2 Decision Model**

Jeff Bezos, Amazon, popularized via shareholder letters

Stage 02.

› **Ladder of Inference**

Chris Argyris, 1970s · popularized by Peter Senge, The Fifth Discipline, 1990

Seven-rung model from raw data to action. Appendix E.

› **Regret Minimization Framework**

Jeff Bezos, 1994; used personally when deciding to found Amazon

Long-horizon distancing technique for identity-laden decisions. Stage 08, Appendix E.

› **Framing Effect**

Kahneman & Tversky, 1981 (Prospect Theory)

Same information, presented differently, produces different decisions. Appendix A.

› **Inconsistency Avoidance Tendency**

Charlie Munger, Poor Charlie's Almanack

Once a position is taken publicly, people defend rather than evaluate it. Appendix A.

› **The Four Blind Spots**

Medium / 'Four Conundrums of Intelligence' — referenced in Ratipcan Uysal, Strategic Thinking, 2024

Out of context, out of mind, lackluster, expired. Stage 06 pre-mortem protocol.

› **Metacognitive Questions**

Ratipcan Uysal, Strategic Thinking: Cognitive Biases & Mental Models, 2024

Self-questioning framework for examining belief formation. Appendix E, Stage 08.

› **OODA Loop — Orient Before You Decide**

Col. John Boyd, U.S. Air Force, 1960s–1990s · via Shane Parrish, Farnam Street, 2018

Observe, Orient, Decide, Act. The Orient step (schwerpunkt), destructive deduction, and tempo. Appendix E.

› **Adversarial Thinking**

Game theory: von Neumann & Morgenstern, Theory of Games and Economic Behavior, 1944

Steelmanning the opponent's best response before committing. Appendix E.

› **Asymmetric Bet Sizing**

J.L. Kelly, A New Interpretation of Information Rate, 1956 · Warren Buffett (20-slot rule) · Nassim Taleb, Antifragile, 2012

Bet sizing, Kelly Criterion, barbell strategy, upside/downside asymmetry. Appendix E.

› **Optionality**

Nassim Taleb, Antifragile, 2012 · real options theory

Decisions that open vs. close future paths; buying vs. selling optionality under uncertainty. Appendix E.

› **Commitment Escalation**

Barry Staw, Knee-Deep in the Big Muddy, Organizational Behavior and Human Performance, 1976

How reversible decisions become irreversible through incremental commitment. Appendix E.

› **10/10/10 — Temporal Perspective**

Suzy Welch, 10-10-10: A Fast and Powerful Way to Get Unstuck, 2009

Three time horizons to neutralize present-bias in high-pressure decisions. Appendix E.

› **Weighted Criteria**

McKinsey / Barbara Minto tradition · Chevallier & Enders, Solvable, 2022

Define criteria and weights before evaluating options to prevent anchoring. Appendix E.

› **Inversion**

Carl Jacobi (19th c.) · Charlie Munger, Poor Charlie's Almanack, 2005 · Jasmine Bina, Concept Bureau, 2025

Think backward: how would I guarantee failure? Engineer its prevention. Appendix E.

› **Local vs. Global Maxima**

Jasmine Bina & JL Rawlence, The Strategist's Playbook, Concept Bureau, 2025

Optimizing the current hill vs. positioning for the next one. Appendix E.