

Creative, Feasible Alternatives

- A common cause of strategic failures comes from choosing too soon: a team finds a workable plan, builds the case for it, and stops looking. The option it never generated is the one it can never choose.
- This is the second link in Decision Quality, after the frame is settled and before any analysis begins. Analysis cannot recover a missing option.
- The discipline is concrete: put three to five genuinely different options on the table, including the status quo, before evaluating anything. A real alternative differs on at least two strategic dimensions, not just price or pace.
- Five practices keep the set open: mandate orthogonal options, use scenarios to build families of them, treat strategy as a portfolio of staged bets, build in structured challenge, and run a room-level prompt to break late narrowing.
- More is not always better. Past roughly five distinct strategies, the returns fall off, and the binding limit is whether the team can still compare them, not the count. For almost every team, though, the real risk is stopping at one.
- Generative AI is useful for generating and restating options, but it cannot set thresholds, weigh stakes, or confer legitimacy. Treat its output as raw material, not the option set.

A common and costly failure in decision-making is settling on a solution too soon. A team finds a workable option, builds the case for it, and stops exploring alternatives. Simon's satisficing explains the pattern: under real limits of time, attention, and information, decision-makers search until an option clears the bar of "good enough", then stop.¹ While this is a rational response to constraints, it quietly sets a ceiling on the outcome because the option a team never generated cannot be chosen.

Creative, feasible alternatives is the second link in Decision Quality, after the frame is settled and before any analysis begins.² Put three to five materially different options on the table, including the status quo, before evaluating anything. The evidence shows that teams rarely do this, and that it pays when they do. This paper covers why capable teams stop at one option and the practices I use to keep the option set open long enough to matter.

The theory in brief (why alternatives determine strategy outcomes)

Satisficing has a mechanism. Cyert and March showed that organizational search is problem-driven: it begins when performance falls short of an aspiration level and ends when a satisfactory option appears.³ Greve confirmed the pattern in the field, with firms changing course when results dip below aspirations and easing off once they recover.⁴ Psychology helps explain why. Under time pressure, people seek **cognitive closure**, seizing an early answer and locking in on it.⁵ These are the defaults of bounded minds and busy organizations, which is why the correction has to come from process rather than exhortation.

The cost shows up in the outcome data. Nutt followed 356 decisions in U.S. and Canadian organizations for two decades and traced failure to managers who impose solutions, limit the search for alternatives, and use power to push the plan through.⁶ Managers developed more than one option in fewer than a fifth of the decisions, and where they did, success rose from 56 to 70 percent.⁶ The most option-rich approach, a search conducted through repeated learning cycles, did not fail in any of the cases where it was applied. However, since this method accounted for only 3 percent of the sample, its statistical significance is limited compared to the broader findings.⁶ Process studies point the same way: teams that follow a rational, information-driven process reach more effective strategic choices,⁷ and Eisenhardt's fast decision-makers in high-velocity markets considered more options at once, not fewer.⁸

Once a plausible path appears, the narrowing compounds. Schwenk cataloged how this happens specifically in strategic decisions: a prior-hypothesis bias recruits evidence for the early favorite, and an anchor sets the reference point against which every later option is judged.⁹ Strategists also reason by analogy more than they admit, and a surface match to another company's situation imports that company's answer along with it.¹⁰ The **inside view** finishes the job: the team builds its case from the plan's own fact base and skips what comparable bets elsewhere would predict. Lovallo and Kahneman's outside view is the corrective, and it doubles as an option generator, because asking what others in this position did returns alternatives alongside base rates.¹¹

This is why alternatives sit second in the Decision Quality chain. The frame settles the question; the option set decides whether the later analysis compares genuinely different futures or rehearses a foregone conclusion. The DQ standard requires options to be truly distinct and genuinely feasible.² For hard-to-reverse decisions, I use this standard as a strict gate: evaluation does not begin until the available options are both meaningfully different and feasible.

From theory to practice (building better option sets)

Each practice below pairs a short reason it works with a description of what good output looks like, so the option set becomes a visible, auditable link rather than a step a team can claim to have done. The frame is already settled when this work starts; if it isn't, go back to it before generating a single option.

A. Mandate *orthogonal* alternatives (not variants)

Why it works.

Requiring three to five genuinely different strategies is the practical counter to satisficing, since a team that has to fill the slots cannot stop at the first plan that clears the bar. Different has to mean something, so the test is that an option varies on at least two strategic dimensions, such as business model, customer segment, capability path, or risk posture. Nutt's data is the warrant: more than one option in fewer than a fifth of decisions, and success up from 56 to 70 percent when teams managed it.⁶

What good looks like.

Three to five options that differ on real dimensions, with the status quo named as one of them. Each carries its decision logic, its load-bearing assumptions, and the kill criteria that would retire it. What you do not have is an approval deck built to walk a single recommendation through.²

B. Use scenarios to design *families* of options

Why it works.

A small set of contrasting futures pushes a team to design strategies for each, which surfaces hedges and staged bets it would otherwise miss and lets you see which option holds up across more than one future.¹² Robust-decision methods formalize the same instinct: prefer an option that performs acceptably across states over one tuned to the single future you find most likely.¹³

What good looks like.

Two or three plausible futures and a simple matrix scoring each option against each, with at least one option that clears a minimum score across all of them rather than winning big in one and failing the rest.

C. Treat strategy as a *portfolio of real options*

Why it works.

Real-options logic breaks one large irreversible bet into staged commitments you can extend or abandon as you learn, which legitimizes a “probe and learn” path next to the core plays.¹⁴ Discovery-driven planning supplies the discipline around it: surface the assumptions each bet depends on, and tie the next tranche of funding to the test that checks them.¹⁵

What good looks like.

A page that separates the few core commitments from several staged bets, each with an explicit learning goal and a trigger to scale up or kill. Funding follows the checkpoints, not the calendar.

D. Build challenge into the search (red teams, premortems)

Why it works.	What good looks like.
Structured dissent works where informal debate does not. In controlled comparisons, groups using dialectical inquiry or devil's advocacy reached higher-quality recommendations and surfaced sounder assumptions than consensus groups. ¹⁶ A premortem does the same job from the other end: assume the plan failed, and the explanations a team generates point to options and safeguards it had not considered. ¹⁷	A short adversarial review of the shortlist by a cross-functional team outside the sponsor's reporting line, timeboxed to about ninety minutes, plus a one-page premortem. The review earns its place when it produces at least two new options or safeguarded variants, such as a staged rollout.

E. Force expansion with simple room tools

Why it works.	What good looks like.
A small intervention can break inside-view lock-in late in the process, when the set has quietly narrowed to a favorite and a straw man. This is a distinct failure from never generating options at all: the team did the work early, then let attrition and sunk commitment collapse a real set back down to one as the deadline approached. The vanishing-options test is the one I rely on. ¹⁸	Before down-selecting, ask the room: assume the current options are all off the table, what would we do instead? Keep any answer that is materially different from what is already on the list, not a faster or cheaper version of it.

Practical limitations (and how to work with them)

- **More options cost time** — so the requirement has to match how reversible the decision is. A one-way door earns the full set of three to five genuinely different options including the status quo; a reversible choice may need only a default and one counterfactual. Give optioning a set slice of the timeline, and stop when another round turns up only variants of what you already have.⁸
- **Generating alternatives threatens the favorite** — the option set is partly a political act, since every option added dilutes the sponsor's preferred plan. Nutt's failures cluster exactly here, in decisions where managers used power to push a single solution rather than weigh several.⁶ The guard is to separate who generates options from who chooses among them, and to settle the decision rights before the work starts.
- **Most “alternatives” are really variants** — a team under pressure will pad the list with the same plan at three speeds and call it range. Hold the orthogonality test from practice A, that a

real option moves on at least two strategic dimensions, and keep the status quo in the set so every option has something honest to be measured against.²

- **You cannot fully analyze every option** — and you do not need to. Screen the full set cheaply against a few decision-relevant criteria, then build detailed cases only for the two or three that survive. That preserves breadth without drowning the team in analysis.
- **More is not always better** — the prescription is “more and more different,” not as many as possible. There is no validated ceiling for strategic options; the real limit is whether the team can still hold them in genuine comparison and whether each new entry is actually distinct. Past roughly five orthogonal strategies and the baseline, the next “option” is usually a variant that adds load without range, and in one study of an unstable industry, a more comprehensive strategic search was associated with worse performance, not better.¹⁹ The asymmetry is worth keeping in mind, though: Nutt's teams developed more than one option in under a fifth of their decisions,⁶ so for most teams the risk is stopping at one, not drowning in twelve.

Generative AI as scaffold (not substitute)

AI excels at generating options. It will produce a dozen first-pass strategies, restate each at the same level of detail so they can be compared side by side, and push toward the variants a team would struggle to reach for on its own.

But the output is not the option set. It is raw material the team still has to cut, combine, and rewrite against its own goals and risk posture before anything earns a place on the list. AI cannot set the thresholds an option has to clear, choose the risk posture, weigh the political stakes, or confer legitimacy on the option the team lands on. Pair the generation step with the reframing work from the framing paper, since a wider frame is what produces genuinely different options in the first place.

Two prompts do most of the work.

- To expand the set: "Generate twelve strategies for this decision that differ on at least two of business model, customer segment, channel, capability base, capital intensity, or time-to-value, and name the assumption each one bets on." Keep the two or three that change how you see the choice.
- To stress the set for sameness, run the vanishing-options test from practice E as a prompt: "Assume our current shortlist is off the table; give five strategies genuinely different from it and from each other."

The guardrails are the same ones that apply elsewhere. The decision owner signs the final set, not the model. Keep a short log of what you asked and what the model drew on. Use an approved workspace and keep proprietary detail out of prompts unless the environment is secured for it. Tag anything factual the model produces as sourced or asserted, with a one-line confidence note,

before it reaches the room. And run one check before you commit: where did the model hand us look-alikes dressed as alternatives?

Bottom line

A strong option set does not guarantee a good decision, but a weak one rules it out, because no amount of later analysis can recover an option the team never put on the table. That is what makes alternatives the link most worth defending: it is cheap relative to the analysis that follows, and the whole chain is capped by it. The discipline comes down to a few habits that hold under pressure — three to five genuinely different options, the status quo named among them, and a deliberate push against the team's first favorite before anyone runs a number. A half-day widening the option set is cheap against the cost of unwinding the wrong choice a year on.

References

1. Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99–118.
2. Spetzler, C., Winter, H., & Meyer, J. (2016). *Decision Quality: Value Creation from Better Business Decisions*. Wiley.
3. Cyert, R. M., & March, J. G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall.
4. Greve, H. R. (1998). Performance, aspirations, and risky organizational change. *Administrative Science Quarterly*, 43(1), 58–86.
5. Kruglanski, A. W., & Webster, D. M. (1996). Motivated closing of the mind: “Seizing” and “freezing.” *Psychological Review*, 103(2), 263–283.
6. Nutt, P. C. (1999). Surprising but true: Half the decisions in organizations fail. *Academy of Management Executive*, 13(4), 75–90.
7. Dean, J. W., & Sharfman, M. P. (1996). Does decision process matter? A study of strategic decision-making effectiveness. *Academy of Management Journal*, 39(2), 368–396.
8. Eisenhardt, K. M. (1989). Making fast strategic decisions in high-velocity environments. *Academy of Management Journal*, 32(3), 543–576.
9. Schwenk, C. R. (1984). Cognitive simplification processes in strategic decision-making. *Strategic Management Journal*, 5(2), 111–128.
10. Gavetti, G., & Rivkin, J. W. (2005). How strategists really think: Tapping the power of analogy. *Harvard Business Review*, 83(4), 54–63.
11. Lovallo, D., & Kahneman, D. (2003). Delusions of success: How optimism undermines executives' decisions. *Harvard Business Review*, 81(7), 56–63.
12. Ramírez, R., & Wilkinson, A. (2016). *Strategic Reframing: The Oxford Scenario Planning Approach*. Oxford University Press.
13. Kalra, N., Hallegatte, S., Lempert, R., Brown, C., Fozzard, A., Gill, S., & Shah, A. (2014). *Agreeing on Robust Decisions: New Processes for Decision Making Under Deep Uncertainty* (Policy Research Working Paper No. 6906). World Bank.
14. Luehrman, T. A. (1998). Strategy as a portfolio of real options. *Harvard Business Review*, 76(5), 89–99.
15. McGrath, R. G., & MacMillan, I. C. (1995). Discovery-driven planning. *Harvard Business Review*, 73(4), 44–54.
16. Schweiger, D. M., Sandberg, W. R., & Ragan, J. W. (1986). Group approaches for improving strategic decision making: A comparative analysis of dialectical inquiry, devil's advocacy, and consensus. *Academy of Management Journal*, 29(1), 51–71.
17. Klein, G. (2007). Performing a project premortem. *Harvard Business Review*, 85(9), 18–19.

18. Heath, C., & Heath, D. (2013). *Decisive: How to Make Better Choices in Life and Work*. Crown Business.
19. Fredrickson, J. W., & Mitchell, T. R. (1984). Strategic decision processes: Comprehensiveness and performance in an industry with an unstable environment. *Academy of Management Journal*, 27(2), 399–423.