Commitment to Action

- Commitment to action is the sixth and last link in Decision Quality, the one that turns a reasoned choice into work people own, resource, and carry through. A decision can clear every upstream bar and still come apart once the room empties.
- A sound choice that no one carries through is worth no more than a poor one. Surveyed companies realize, on average, only about two-thirds of the value their strategies promise, and much of the shortfall comes after the decision is made.
- Commitment fails at three levels: the individual who means to act and does not, the organization where discussion stands in for action, and the change itself, which slips back unless it is locked in before attention moves on.
- Six moves make commitment hold: assign decision rights, pre-commit structurally, sequence the change and clear the barriers, set exit triggers, log and review on a cadence, and hold a disagree-and-commit norm. Together, they make acting on the choice cost the organization less than letting it lapse.
- Match the apparatus to the stakes and to the reversibility of the call. A one-way door with large consequences earns the full sequence; a reversible bet needs little more than a named owner and a date.
- AI will draft every artifact of commitment in minutes, but the artifacts are not the commitment. A model cannot make anyone own the call, pull a trigger when the sponsor wants to press on, or commit a colleague who disagrees.

A well-made strategy can still come to nothing, and often does. In a Marakon Associates survey of senior executives at 197 large companies worldwide, Mankins and Steele found that firms realize, on average, about 63 percent of the financial value their strategies promise, with the rest lost to breakdowns spread across planning and execution.[1] Some of that gap is set upstream, in the frame, the options, or the reasoning. But a decision can clear every one of those bars and still come apart once the room empties: ownership blurs, the work competes with everything already on the calendar, and within a quarter, the organization has drifted back to what it did before.

Commitment to action is the sixth and last link in Decision Quality, and the one that turns a reasoned choice into work people own, resource, and carry through.[2] It inherits the upstream work, the frame, the options, the information, the priorities, and the reasoning, and it is the link a capable team is most likely to assume will handle itself. This paper explains why sound choices stall after they are made and the moves I use to make commitment hold.

The theory in brief (why commitment determines outcomes)

Decision Quality treats a decision as unfinished until the people who have to act on it are committed to doing so. Commitment to action is one of the six requirements a sound decision must meet; it is not merely a step that follows the decision.[2] The difficulty is that choosing and doing are different kinds of work, and the second does not follow automatically from the first. Even when a team arrives at the right answer through careful reasoning, their efforts can still amount to nothing, since the factors that undermine a decision often exist beyond the logic that shaped it. These factors operate at three levels: the individual, the organization, and the evolution of the change.

Considering the individual level first. Even someone who fully intends to act often does not. Gollwitzer and Sheeran's meta-analysis found that intention alone is a weak predictor of behavior. One of the best-supported ways to bridge this gap is through an **implementation intention**, an if-then plan that links a named cue and the response to it ("if a decision passes its target date without a result, then it goes on the next leadership agenda as a stop-or-continue call").[3] Pre-committing the response to a specific trigger does two things: it makes the cue easier to notice, so the moment is less likely to slip by, and it lets the action begin without a fresh decision once the cue arrives.

At the organization level, when firms know the right course of action, they often fail to execute because discussion substitutes for action, and the fear of making mistakes prevents people from acting on their knowledge. Pfeffer and Sutton termed this the **knowing-doing gap**.[4] Decisions are made, debated, and documented, but the discussion can create the illusion of progress, and no one notices that the actual implementation never starts.

Regarding the change itself, even when it initially succeeds, it can still come apart over time. Rogers describes a confirmation stage that follows the decision and the first implementation, in which the new practice is either integrated and reinforced or quietly abandoned. Until it is, the change remains reversible.[5] This is why initiatives that look successful at launch often regress a few quarters later: the early momentum was real, but it was not integrated into the day-to-day operations. Kotter warns management against declaring victory too soon and failing to embed the change into the organization's routines.[6]

The consensus in the literature is to embed follow-through within organizational structures, making it easier to act on a decision than to let it fade. Establishing clear decision rights, naming who decides and who performs, prevents a settled choice from being silently reopened.[7] Structural pre-commitment, such as public declarations or reallocating budgets, makes action the default rather than an act of will.[8] And reviewing decisions on a fixed cadence, the way Mankins and Steele argue planning should be organized around decisions rather than the calendar, keeps a choice alive as attention moves to the next thing.[9] The rest of this paper translates these principles into actionable steps for teams.

From theory to practice (six moves that make commitment stick)

Each move below pairs a short explanation of why it works with what good output looks like, so commitment becomes a visible, auditable link rather than a step a team claims to have taken. By the time this work starts, the choice is made; the job now is to keep it from coming undone.

A. Assign decision rights & owners (RAPID)

Why it works.	What good looks like.
A settled choice gets quietly reopened when no one is accountable for carrying it out. RAPID names the roles, in particular the single person who decides and the people who perform, so the decision has an owner and stops being relitigated.[7] It matters most where authority is shared, since that is where accountability otherwise diffuses.	The memo names the roles as people, not titles: who **R**ecommends, whose **A**greement is required, who will **P**erform the work, who gives **I**nput, and the single person who **D**ecides. One decider, not a committee. Each performer carries a first milestone, a date, and a measure. The roles are visible to everyone the decision touches, so it is clear who can act and who cannot reopen the call.

B. Pre-commit, structurally

Why it works.	What good looks like.
Intention on its own is a weak predictor of action. An **implementation intention**, a specific if-then plan fixing when and where the action happens, is what reliably closes the gap.[3] A plan still has to survive competing demands, so the stronger move is to change the default: move the budget, block the calendar, change the system setting, so that acting on the decision becomes the path of least resistance and lapsing takes effort.[8]	For each committed action, an if-then plan naming the trigger and the response ("when the quarter opens, the reallocated headcount transfers"), plus at least one structural lock already in place: budget moved, the recurring review on the calendar, the default changed in the system. The test is whether doing nothing now requires someone to actively undo something, rather than letting attention drift.

C. Sequence the change and clear the barriers

Why it works.	What good looks like.
A sound decision still stalls when the things that must be true first, the skills, the systems, the policy changes, lag the main effort, so the work hits walls that were predictable.[6] Removing an obstacle often does more than adding a resource, since the knowing-doing gap is sustained as much by what blocks action as by what is missing.[4] Early visible results matter for a separate reason: adoption stays reversible until the new practice is routinized and reinforced, and a first win is what carries it into that confirmation stage rather than letting it regress.[5]	A one-page sequence of enabling moves, each with an owner and a date, ordered so prerequisites land before the main effort. A short list of barriers to remove, policies, system limits, organizational boundaries, tracked as deliberately as the things being built. One or two near-term wins chosen because they are visible and achievable, scheduled early enough to reinforce the change before attention moves on.

D. Set exit triggers and stop-or-scale rules

Why it works.	What good looks like.
Commitment to a decision is not commitment to it forever, and the two are easy to confuse. Staw showed that people escalate their commitment to a failing course of action, putting more behind it because they already have, and the pull is strongest on the person who made the original call.[10] The structural locks that keep a good choice from lapsing are part of what a trigger has to override, since the same defaults also raise the cost of abandoning a bad one. An exit trigger set before the outcome is known takes the stop-or-continue judgment out of the hands of those most invested in continuing. A premortem does the front-end work: assume the decision has failed, and the explanations the team generates become the indicators worth watching.[11]	For each committed bet, a stop-or-scale rule fixed in advance: the observable indicator, the level that trips it, and the action it triggers, whether that is scale, pause, or stop. A premortem run before commitment, its leading failure modes converted into those indicators. Reversible calls have the decider delegated down with a reversal trigger attached; one-way doors are gated in stages instead.[12] A named owner for each review, distinct from the sponsor where the structure allows, so the trigger is pulled rather than explained away.

E. Log the decision and review it on a cadence

Why it works.	What good looks like.
Organizations have short memories, and a settled decision fades from view as attention moves to the next one. Left alone, it is either quietly dropped or silently reopened, neither of which anyone actually decides. A standing record and a fixed review keep the choice in front of the people accountable for it and make reopening a deliberate act that has to be justified.[7] Mankins and Steele argue that planning should be organized around decisions and their follow-through rather than the annual calendar, which is what a decision cadence puts into practice.[9]	A living decision log (Exhibit — Decision Log), one row per decision at the moment of commitment: title, date, RAPID roles, status, next milestone, review date, and a link to the memo. A short standing item at the top of the leadership meeting that walks the live decisions, their status, and any barrier holding them up. Reopening a settled decision requires new evidence or a changed assumption, stated as such, not a re-run of the original argument. The original reasoning is archived with the entry, so a later challenge meets the record rather than a blank page.

F. Build the norm: disagree and commit

Why it works.	What good looks like.
A decision that part of the room argued against will hold only if those people commit to it anyway, and whether they do turns on how the decision was reached. Tyler's work on procedural justice found that people accept outcomes they dislike when the process was fair, when they were heard and the reasoning was consistent and explained.[13] In strategic decision-making teams specifically, Korsgaard, Schweiger, and Sapienza found that members whose input was genuinely considered showed greater commitment to the decision, attachment to the team, and trust in its leader.[14] Kim and Mauborgne found the same pattern in strategy practice: managers committed to decisions that went against them when the process engaged them, explained the thinking, and made expectations clear.[15] Where dissent carries a cost it does not disappear; it goes quiet and returns as passive resistance.	A stated norm that separates the two phases: challenge the decision hard before it is made, commit to it once it is. Dissent recorded rather than suppressed, so committing does not require anyone to pretend they agreed. The reasoning explained to the people who have to execute it, including why the alternatives were set aside. Undermining a made decision treated as a process violation, while bringing genuinely new evidence to reopen it is not, and the log makes that difference legible.

Exhibit — Decision Log

Log major decisions in real time. Each row is a decision at the moment of commitment. Keep the 'Next review' date current; reopen only by exception. Use short links to the memo and exhibits.

DECISION LOG

Decision log (last 12 months)

ID / Title	Date	Door	Decider (D)	Performer(s)	R/A/I summary	Rationale link	Next review	Status	Reopen rule
DL-23-017 / EMEA pricing guardrails	2025-06-03	Type-2	S. Patel	RevOps + EMEA Sales	R: PM; A: Legal; I: FP&A	memo/EMEA-pricing-017	2025-07-15	On track	SVP Sales may reopen

Fields: Door = Type-1 (one-way) or Type-2 (two-way). R/A/I = Recommend / Agree / Input. 'Reopen rule' names who can reopen and how. Keep links short (e.g., wiki IDs).

Practical limitations (and how to work with them)

- **RAPID cuts across the real org chart** — In matrixed or dual-reporting structures, assigning a decider can create conflict if that person doesn't control the budget or line authority. Forcing the role onto someone without real power leads to decisions no one can enforce. The decider should be the person who both controls the necessary resources and is accountable for the outcome. If these responsibilities are split, make that division clear instead of masking it with a shared title. Identify the few individuals authorized to reopen the decision, and remember that RAPID should map a single decision, not serve as a permanent organizational chart.[7]
- **An anchored change still fades once attention moves** — Momentum peaks during kickoff but quickly wanes. If the change isn't embedded into the organization's day-to-day operations, it will slip back once the initial enthusiasm subsides and the confirmation stage is never reached.[5] The main pitfall is declaring victory too early at the first sign of progress.[6] Instead, establish the anchors, such as process changes, metrics, incentives, and the role that now owns the new way of working, before calling the initiative complete, and keep the win open until these elements are firmly in place.
- **Structure cannot make people act where acting is punished** — Decision rights and if-then plans remain unused when the risk of a visible mistake outweighs the comfort of doing nothing, a dynamic Pfeffer and Sutton identified as fueling the knowing-doing gap.[4] These structures only work when taking action feels safe enough. Make it easy to take small, reversible steps, evaluate decisions based on their reasoning rather than the luck of the outcome, and eliminate penalties that encourage hesitation.[12]
- **Competing priorities quietly outcompete the committed action** — A well-owned decision still loses to the day job, because the day job has deadlines and the decision has good intentions. The guard is to give the committed work the same standing as everything else that gets measured: convert the top enabling moves into the goals the organization actually tracks, and tie funding and review to the decision's milestones rather than to the planning calendar.[9]
- **A review can run green while the behavior never changes** — Status reports reflect what is reported, not necessarily what is happening. A decision can appear on track in every review while the actual work quietly reverts to the old way, because it's easier to report a change than to make one. The log and the review catch a stalled decision, not a hollow one. Guard against this by reviewing tangible evidence of new behaviors rather than relying solely on status updates, and by periodically inspecting the work itself rather than just the dashboard, ensuring progress is measured by real change.
- **The full apparatus can cost more than the decision is worth** — RAPID, a log entry, exit triggers, and a standing review are real overhead. For small, reversible decisions, this overhead can slow the organization without adding value. Align the process with the level of risk and reversibility: use the full sequence for high-stakes, one-way decisions, but for reversible bets, a named owner and target date are often sufficient.[12]

Generative AI as scaffold (not substitute)

AI can draft every artifact of commitment in the time it takes to read the memo. What it cannot produce is the commitment itself, because that lives in the people named in the artifacts and nowhere else.

Take them one at a time. AI will propose a RAPID map with named candidates, but it cannot make any of those people own the call. It will turn each action into an if-then plan, but the plan commits no one until the person who has to act agrees to it. Ask it for a premortem, and it will list failure modes to seed the exit triggers, but it cannot pull a trigger when the indicator trips and the sponsor wants to press on. It will keep the decision log and flag what is due for review, but a log records commitments; it does not create them. It can record a dissent, but it cannot win over the person who lodged it.

So the drafting is worth handing off. It clears the mechanical work off the team's plate and puts their attention where it has to be, on the judgments the artifacts can't make: who should own the call, whether the triggers are the right ones, and who still needs to be brought to commit.

A few prompts do most of the work. Adapt the wording to your decision and your data.

- To draft the package: "From this decision memo, draft the commitment package: a RAPID map with named candidates, an if-then plan for each committed action, the exit triggers, and one or two near-term wins. Flag any action with no clear performer." Then put real names against the roles yourself; the model can suggest, not assign.
- To build the change-path: "From these policy and process documents, list the enabling moves this decision depends on and the barriers most likely to block it, in the order they have to be cleared." Verify each dependency is real before you sequence the work, since a model will assert a constraint as readily as it finds one.
- To prep the cadence review: "From this decision log, list the items due for review in the next two weeks with owner, status, and next milestone, and flag any missing a named decider or performer." Use it to prepare the review, not to run it.

The guardrails are the same across the series. The decision owner signs the role map, the triggers, and the change path, 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. Require a source for any claim the model makes about how the organization actually runs, since it will state an assumed constraint as confidently as a real one. And run the one check that matters most for this link: where did the document the model produced stand in for the commitment it only describes? A clean role map and a populated log are the easiest things in the world to mistake for the act of committing, which is the knowing-doing gap in a new form.[4]

Bottom line

By the time you reach this stage, the preferred choice has been identified, but the decision is not yet complete. Commitment determines whether the organization is genuinely prepared to own, resource, and act on that choice. A well-reasoned recommendation that no one funds or executes remains potential value rather than realized value. The key is to embed follow-through in the decision itself rather than treating it as an afterthought.

Of course, none of this ensures success; the outcome was never fully within the team's control. What these practices safeguard is the work already done, so a decision reached at real cost isn't quietly lost in the aftermath.

References

1. Mankins, M. C., & Steele, R. (2005). Turning great strategy into great performance. *Harvard Business Review*, 83(7), 64–72.
2. Spetzler, C., Winter, H., & Meyer, J. (2016). *Decision Quality: Value Creation from Better Business Decisions*. Wiley.
3. Gollwitzer, P. M., & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. *Advances in Experimental Social Psychology*, 38, 69–119.
4. Pfeffer, J., & Sutton, R. I. (2000). *The Knowing-Doing Gap: How Smart Companies Turn Knowledge into Action*. Harvard Business School Press.
5. Rogers, E. M. (2003). *Diffusion of Innovations* (5th ed.). Free Press.
6. Kotter, J. P. (1995). Leading change: Why transformation efforts fail. *Harvard Business Review*, 73(2), 59–67.
7. Rogers, P., & Blenko, M. (2006). Who has the D? How clear decision roles enhance organizational performance. *Harvard Business Review*, 84(1), 52–61.
8. Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale University Press.
9. Mankins, M. C., & Steele, R. (2006). Stop making plans; start making decisions. *Harvard Business Review*, 84(1), 76–84.
10. Staw, B. M. (1976). Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance*, 16(1), 27–44.
11. Klein, G. (2007). Performing a project premortem. *Harvard Business Review*, 85(9), 18–19.
12. Bezos, J. (2015). *2015 Letter to Shareholders*. Amazon.com, Inc.
13. Tyler, T. R. (2006). Psychological perspectives on legitimacy and legitimation. *Annual Review of Psychology*, 57, 375–400.
14. Korsgaard, M. A., Schweiger, D. M., & Sapienza, H. J. (1995). Building commitment, attachment, and trust in strategic decision-making teams: The role of procedural justice. *Academy of Management Journal*, 38(1), 60–84.
15. Kim, W. C., & Mauborgne, R. (1997). Fair process: Managing in the knowledge economy. *Harvard Business Review*, 75(4), 65–75.

BREAKOUT
GROWTH PARTNERS

Kerr Manson
(847) 602-4275
kerr.manson@breakout-growth.com
www.breakout-growth.com