

RESEARCH & PERSPECTIVES

The thinking behind the thesis

A Framework for Making Good Decisions

Process, Judgment, and Commitment

Research & Perspectives

The Decision-Making Guide for Investments

2026

Abstract

This paper proposes a practical framework for making high-quality decisions, drawing on three bodies of research that have each stood up to decades of scrutiny: the decision analysis tradition, the psychology of judgment, and field evidence on what actually helps. The method centres on six elements of decision quality — an appropriate frame, creative alternatives, relevant and reliable information, clear values and tradeoffs, sound reasoning, and commitment to action — combined with four bias-reduction practices: independence before discussion, mediating assessments, the premortem, and process interrogation. A tiering rule matches the depth of analysis to the stakes involved, and a system of records, reviews, and roles makes the discipline self-correcting. The goal is not to eliminate uncertainty but to ensure that what can be controlled — the quality of the process — is controlled, so that the odds tilt reliably toward better outcomes over time.

Keywords: *decision quality; decision analysis; cognitive bias; premortem; mediating assessments; six elements; noise reduction; investment process; judgment under uncertainty*

1. Introduction: Why Good People Make Bad Decisions

Most of us judge decisions by how they turn out. A deal that made money was a good decision; a deal that lost money was a bad one. This feels like common sense, and it is almost always wrong. Outcomes are the joint product of decision quality and luck, and luck does not care how carefully you thought. A poker player who goes all-in with a weak hand and wins has still played badly. An investor who did everything right and got blindsided by a pandemic has still invested well. Ronald Howard, who founded the modern discipline of decision analysis at Stanford in the 1960s, put the point simply: the quality of a decision can only be judged at the moment it is made, by the quality of the process that produced it, never by what happened afterwards (Howard, 1988).

This distinction is not a philosopher's quibble. It is the foundation of every serious method for deciding well, because it redirects our attention from the one thing we cannot control, the outcome, to the one thing we can, the process. And the evidence that process matters is unusually strong. When Dan Lovallo and Olivier Sibony studied 1,048 major business decisions over five years, comparing the analytical work behind each decision with the process used to reach it, they found that process mattered roughly six times as much as analysis in explaining which decisions created value (Lovallo & Sibony, 2010). Superb spreadsheets fed into a broken process produce broken decisions. A sound process, on the other hand, tends to catch weak analysis before it does harm, because a sound process asks where the numbers came from and what would have to be true for them to hold.

The framework set out in this document rests on that finding. It draws on three bodies of work that have each stood up to decades of scrutiny: the decision analysis tradition begun by Howard and codified by Carl Spetzler and his colleagues as the six elements of decision quality (Spetzler et al., 2016); the behavioural research of Daniel Kahneman and his collaborators on bias and noise in human judgment (Kahneman, 2011; Kahneman et al., 2021); and the practical toolkit that has grown up around both, including Gary Klein's premortem and the structured evaluation methods known collectively as multi-criteria decision analysis (Klein, 2007; Belton & Stewart, 2002). None of these alone is sufficient. Together they form something close to a complete method: an architecture that tells you what a good decision must contain, a discipline that protects your judgment while you build it, and an engine for comparing the options once they are on the table.

A word of reassurance before we begin. None of this requires a committee, a consultant, or a month of meetings. The full method is reserved for the decisions that warrant it. Most decisions do not, and knowing the difference is itself part of the framework.

2. What the Research Actually Tells Us

Three streams of research converge on the same conclusion from different directions, and it is worth seeing each one clearly before assembling them into a working method.

The first stream is decision analysis, an applied branch of operations research and economics. Its central contribution is structural. A decision, it argues, is not a mysterious act of will but a thing with parts, and each part can be examined and improved. Spetzler, Winter, and Meyer (2016), building on fifty years of consulting practice at the Strategic Decisions Group, identified six such parts: the frame, the alternatives, the information, the values, the reasoning, and the commitment to act. Their most useful observation is that these six behave like links in a chain. A decision is only as strong as its weakest element. Brilliant analysis of the wrong question is worthless. A perfectly framed problem with only one alternative under consideration is not a decision at all; it is a foregone conclusion wearing the costume of one.

The second stream is the psychology of judgment, which for fifty years has documented the systematic ways human thinking goes wrong. Kahneman and Tversky's work on heuristics showed that our errors are not random; they lean, predictably, in particular directions. We anchor on the first number we hear. We seek evidence that confirms what we already believe. We are overconfident about our forecasts and overly attached to what we already own. More recently, Kahneman, Sibony, and Sunstein (2021) added a second and less familiar category of error: noise, the sheer variability of judgment. Two experienced underwriters given the same file will price it differently, often wildly so, and the same underwriter will price it differently on Monday than on Thursday. Bias is the average error; noise is the scatter. Both degrade decisions, and the remedies for them are largely procedural rather than heroic. You cannot will yourself into objectivity, but you can design a process that does not depend on it.

The third stream is the field evidence on what actually helps. Here the news is encouraging. Klein's premortem, in which a team imagines a decision has already failed and writes down why, measurably increases the number of risks identified and does so in an hour (Klein, 2007). Kahneman, Lovallo, and Sibony (2011) distilled the debiasing literature into a twelve-question checklist for reviewing recommendations before approving them, aimed not at the content of a proposal but at the process and incentives that produced it. Tetlock's forecasting tournaments showed that ordinary people who break problems into parts, use outside reference classes, and update their views in small increments consistently beat credentialed experts who do none of these things (Tetlock & Gardner, 2015). The pattern across all of it is the same. Good judgment is less a talent than a set of habits, and the habits can be written down.

3. The Architecture: Six Elements of a Good Decision

The six elements of decision quality are the spine of this framework. They are best understood not as steps in a sequence but as tests. Before committing to any significant decision, each element should be able to pass its test, and the honest question is always which link is currently the weakest.

The first element is an appropriate frame. Framing is the act of deciding what, exactly, you are deciding. It sounds trivial and is the most commonly botched of the six. The frame fixes the purpose of

the decision, its scope, and the perspective from which it is viewed. Consider a family office reviewing an underperforming holding. Framed narrowly, the question is whether to sell. Framed properly, it may be whether this asset class still belongs in the portfolio at all, or whether the underperformance is information about the thesis rather than the asset. A useful test: state the decision in one sentence, then ask whether someone you respect would say you are answering the wrong question. Too narrow a frame hands you a tidy answer to an irrelevant problem. Too wide a frame turns a decision into a philosophy seminar. The skill lies in choosing deliberately rather than inheriting the frame from whoever raised the issue first.

The second element is a set of creative, genuinely different alternatives. You cannot choose an option you never generated, and the quality of a decision is capped by the best alternative on the table. The most frequent failure here is the binary trap: do the deal or don't, hire or don't, sell or hold. Chip and Dan Heath (2013) call this a spotlight problem; the light sits on one option while the rest of the stage stays dark. The remedy is mechanical and works: require at least three genuinely distinct alternatives before any evaluation begins, and include where honest a hybrid, a deferral with a defined trigger, and the vanishing-options test, which asks what you would do if the current favourite were forbidden. Alternatives should be different in kind, not variations of one idea wearing three hats.

The third element is relevant and reliable information. Notice both adjectives. Information is relevant when it bears on the decision at hand, and reliable when it would survive hostile cross-examination. Much of what circulates in any deliberation is neither: it is vivid, recent, and secondhand. Two practices raise the standard sharply. The first is the outside view. Before trusting your inside forecast of this project, this fund, this renovation, look at the base rate of similar cases. Most projects of this type run over budget by a knowable percentage; most acquisitions fail to earn their premium; most first-time authors sell fewer than a thousand copies. Starting from the base rate and adjusting is consistently more accurate than building up from the particulars (Kahneman, 2011). The second practice is calibrating uncertainty honestly. Replace point estimates with ranges, and ask for the ten percent downside case in writing. A single number is a forecast pretending to be a fact.

The fourth element is clear values and tradeoffs. A decision without explicit values is a boat without a rudder; it will go somewhere, but not on purpose. Values here means the criteria by which the alternatives will be judged and the relative weight each criterion carries. Return, risk, liquidity, time, reputation, family harmony, the demands of conscience: these rarely point the same way, and refusing to weigh them does not avoid the tradeoff, it merely makes the tradeoff silently and by default. Ralph Keeney (1992) argued that values, not alternatives, should be the starting point of hard decisions, and there is real force in that. Write the criteria down before evaluating anything, because criteria invented after the fact have a way of arranging themselves conveniently behind the option someone already prefers.

The fifth element is sound reasoning. Once the frame is set, the alternatives are generated, the information is gathered, and the values are explicit, reasoning is the act of putting them together

honestly. For most decisions this needs nothing more exotic than a clear head and a written argument. For complex ones it may mean a weighted scoring model, a decision tree, or a formal expected-value calculation. The form matters less than one property: the reasoning must be visible. An argument that exists only in someone's head cannot be checked, and an argument written down has a wonderful tendency to expose its own gaps to its author before anyone else sees it.

The sixth element is commitment to action. A decision that does not move anyone is not a decision; it is an opinion with minutes. This element covers two distinct failures. The first is the analysis that ends in a drawer, endlessly refined and never executed. The second is subtler: execution by people who were never persuaded, who comply outwardly while quietly starving the decision of effort. The best protection against both is involving the people who must act early enough that their objections improve the decision instead of sabotaging the implementation. Commitment is built during the process, not requested after it.

4. The Discipline: Keeping Bias and Noise Out

The six elements describe what a good decision contains. They say nothing about the minds doing the containing, and those minds are the problem. Every element can be quietly corrupted: frames inherited from the loudest voice in the room, alternatives generated to make a favourite look good, information selected to confirm, values weighted after the fact, reasoning bent toward the desired conclusion, commitment manufactured through pressure. Awareness of bias, the research is blunt on this point, does very little to prevent it (Kahneman et al., 2011). Knowing that anchoring exists does not stop you anchoring, any more than knowing about optical illusions makes the lines look equal. What works is not self-knowledge but procedure, what Kahneman, Sibony, and Sunstein (2021) call decision hygiene, on the analogy of hand-washing: unglamorous, preventive, and effective against pathogens you cannot see.

Four practices carry most of the weight, and none takes long.

First, independence before discussion. The moment a group starts talking, its members start influencing one another, and the information held by quieter or more junior people begins to evaporate. The fix is to collect every participant's view in writing before the discussion opens. Estimates, concerns, preferred options, all committed privately first. This single habit preserves more independent information than any facilitation technique, and it costs ten minutes. In a family office or a small board the temptation to skip it is strong precisely because the group is intimate; intimacy is exactly what makes early convergence so fast and so invisible.

Second, assess before you judge. Kahneman, Lovallo, and Sibony's mediating assessments protocol asks the group to break the decision into its key dimensions, score each dimension separately on the evidence, and postpone any overall verdict until every dimension has been scored. The point is to stop

the halo effect, in which an early good impression of one aspect quietly inflates the scores of all the others. An acquisition target with a charismatic founder somehow acquires strong financials, a defensible market, and manageable legal risk, all before anyone has checked. Keeping the assessments separate, and looking at them side by side only at the end, is the closest thing judgment research has produced to a general-purpose safeguard (Kahneman et al., 2021).

Third, run a premortem. When a decision is close to being made, Klein's exercise takes one hour. The group is told to imagine that it is two years on and the decision has failed, plainly and undeniably. Each person independently writes the history of that failure. The framing matters: not 'what could go wrong,' which invites polite generalities, but 'it went wrong, explain why,' which licenses the doubts people have been sitting on. The premortem is the cheapest insurance in the entire toolkit, and it has a second benefit: it makes dissent a task rather than a betrayal, which changes what people are willing to say (Klein, 2007).

Fourth, interrogate the process, not just the proposal. Before approving any significant recommendation, including your own, run it through a short version of the Kahneman, Lovaglio, and Sibony checklist. Is there any reason to suspect the recommenders are attached to their conclusion? Were credible alternatives seriously examined, or ceremonially listed? If this same proposal came from someone you disliked, would the evidence still persuade you? Where is the base rate, and why should this case beat it? Is the worst case the real worst case, or the worst case that remains comfortable to present? A proposal that survives these questions has earned its approval. One that resents them has answered them.

A note on intuition, because the discipline described here can sound like a war against it. It is not. Klein's own research on firefighters, nurses, and chess players shows that intuition is genuinely expert in domains that offer regular patterns and fast, unambiguous feedback (Klein, 1998). The trouble is that markets, strategy, and most of the decisions this framework exists for are not such domains; feedback arrives late, entangled with luck, and easy to misread. The working rule is this: let intuition propose, and let process dispose. A strong gut feeling is a hypothesis worth taking seriously, and never a verdict.

5. The Engine: Comparing the Alternatives

With the elements in place and the hygiene observed, there remains the mechanical question of how to compare the options. The right tool depends on the structure of the problem, and it is worth resisting the belief that more machinery means more rigour. Machinery amplifies whatever judgment is fed into it, including bad judgment, with impressive precision.

For decisions with a handful of alternatives and a handful of criteria, which is most of them, a simple weighted scoring table is enough. List the criteria agreed in the values stage, weight them so the weights sum to one hundred, score each alternative against each criterion, and multiply through. The

output is not the answer; it is a structured summary of your own stated beliefs, and its chief virtue is what it exposes. When the table says one thing and your stomach says another, something has been left out of the table or the stomach is wrong, and finding out which is precisely the conversation worth having. Sensitivity checking takes minutes: change the weights within reasonable bounds and see whether the ranking flips. A conclusion that survives every plausible weighting is sturdy. One that flips on small changes is telling you the decision is genuinely close, which is also worth knowing.

For decisions dominated by uncertainty rather than by competing criteria, the natural tool is the decision tree and expected value. Lay out the choice points, the chance events with honest probabilities, and the outcomes with honest magnitudes, then fold the tree back. The discipline here is in the probabilities. Use ranges, consult base rates, and remember Tetlock's forecasters, whose accuracy came from granular, frequently updated probability estimates rather than from confident stories (Tetlock & Gardner, 2015). Expected value has a known blind spot: it is indifferent to ruin. A ten percent chance of losing everything is not offset by a ninety percent chance of a pleasant gain, not for a family, and the framework therefore adopts a hard rule that predates all of this literature: no expected value, however attractive, justifies exposure to unacceptable loss. Survival constraints are screened first, and only the alternatives that pass are compared.

For genuinely complex evaluations with many stakeholders and many criteria, formal multi-criteria decision analysis offers well-tested procedures for structuring criteria, eliciting weights, and aggregating scores (Belton & Stewart, 2002). In an institutional setting these methods earn their overhead by making the reasoning transparent and contestable. In a personal or family setting they are usually more apparatus than the decision needs, and the simple weighted table, honestly filled in, captures most of the benefit.

One further variable governs how much of this engine to deploy: reversibility. A reversible decision is an experiment, and the fastest way to learn from an experiment is to run it. An irreversible decision is a commitment, and commitments deserve the full treatment. Asking 'how expensive is it to be wrong, and how quickly would we know' is often more valuable than any amount of additional analysis, because the answer frequently converts an agonising decision into a cheap trial.

6. Matching the Effort to the Stakes

A framework that treats every decision like a merger will be abandoned within a month, and deservedly. The final structural piece is a tiering rule that matches process to consequence.

Tier one covers routine and reversible decisions: most purchases, most scheduling, most operational choices. The full framework would be absurd here. The only requirements are a clear frame in one sentence, a moment's check that an obvious alternative is not being ignored, and speed. Deciding quickly and moving on is the correct method, and time recovered here funds the tiers above.

Tier two covers significant but recoverable decisions: a material investment within an established strategy, a key hire, a major commitment of time. Here the framework runs in its light form. The decision is framed in writing, at least three real alternatives are generated, the relevant base rate is looked up, criteria and weights are set before evaluation, a scoring table or short written argument compares the options, and one hygiene practice, usually the premortem or the checklist, is applied. Elapsed time: a few hours, perhaps spread over days to let first impressions cool. The written record matters as much as the analysis, for reasons the next section explains.

Tier three covers rare, high-stakes, hard-to-reverse decisions: selling a business, restructuring a portfolio's foundations, a family commitment that shapes a decade. Here everything applies. Full framing with deliberate consideration of wider and narrower versions of the question. Alternatives generated to a quota, with the vanishing-options test. Independent written views collected before any group discussion. Mediating assessments scored separately. The outside view documented. A full premortem. The process checklist applied without mercy. And a waiting period between the provisional conclusion and the final commitment, because a conclusion that still looks right after two weeks of living with it is a different creature from one that merely survived a persuasive meeting.

The tiering rule itself takes ten seconds to apply. What is at stake, and can it be undone? Those two questions sort almost every decision correctly, and the discipline of asking them is the habit on which the rest of the framework hangs.

7. Making It Stick: Records, Review, and Roles

Frameworks fail in the implementation, not the design, and the failure is always the same: the method is used enthusiastically for a month and then quietly abandoned when it becomes inconvenient, which is to say, when it is needed. Three practices give it staying power.

The first is the decision log. For every tier two and tier three decision, a short record is kept at the moment of commitment: the frame, the alternatives considered, the key assumptions and probabilities, the reasons, and what would have to happen for the decision to be revisited. Half a page suffices. The log serves two masters. It disciplines the present, because writing for the record raises the standard of the reasoning. And it educates the future, because memory is the least reliable instrument we own. Hindsight quietly rewrites what we knew and when we knew it, and without a log, every past decision becomes evidence for whatever we currently believe.

The second is the review. Once or twice a year, the log is reread against results, asking two separate questions that must never be merged. Was the outcome good? And was the process good? Four combinations exist, and the instructive ones are the off-diagonal cases: good process followed by a bad outcome, which deserves study but not regret, and bad process followed by a good outcome, which deserves the sternest scrutiny of all, because rewarded carelessness is how bad habits become

convictions. The review is also where calibration lives. If the confident forecasts in the log keep failing at a higher rate than their stated probabilities, the forecasts were not honest, and the next log entry should know it.

The third is a clear structure of roles, even in a small setting. Every significant decision should have a named decision-maker, a defined set of advisors whose views are collected independently, and, for tier three, an explicitly appointed dissenter whose task is to argue the strongest case against the emerging consensus. Assigning dissent as a role does something subtle and valuable: it separates the argument from the arguer, so that the case against can be made at full strength without anyone spending personal capital to make it. Where a decision touches family, the same logic argues for involving affected members early, in the alternatives stage rather than the announcement stage, both because they hold information and because commitment to action, the sixth element, is built there or not at all.

8. Closing: Decision-Making as a Practice

There is an old Stoic distinction, and it belongs at the end of this document because the entire framework is an elaboration of it. Epictetus opens the *Enchiridion* by dividing the world into what is up to us and what is not. Our judgments, our aims, our efforts are ours; outcomes, in the end, are not. The framework asks nothing more than that this line be taken seriously. Frame the question deliberately. Generate real alternatives. Weigh honest information against explicit values, with reasoning that can be shown to another person. Protect the whole process from the predictable failures of your own mind, commit fully, record faithfully, and review without flinching. Then accept the outcome, whatever it is, as the part that was never yours to command.

Practised this way, decision-making stops being a series of anxious verdicts and becomes something closer to a craft: a repertoire of habits that improve with honest use, that can be taught, and that compound. No single decision will reveal whether the framework is working. A decade of them will.

References

- Belton, V., & Stewart, T. J. (2002). *Multiple criteria decision analysis: An integrated approach*. Kluwer Academic Publishers.
- Heath, C., & Heath, D. (2013). *Decisive: How to make better choices in life and work*. Crown Business.
- Howard, R. A. (1988). Decision analysis: Practice and promise. *Management Science*, 34(6), 679–695.
<https://doi.org/10.1287/mnsc.34.6.679>
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kahneman, D., Lovallo, D., & Sibony, O. (2011). Before you make that big decision. *Harvard Business Review*, 89(6), 50–60.
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment*. Little, Brown Spark.
- Keeney, R. L. (1992). *Value-focused thinking: A path to creative decisionmaking*. Harvard University Press.
- Klein, G. (1998). *Sources of power: How people make decisions*. MIT Press.
- Klein, G. (2007). Performing a project premortem. *Harvard Business Review*, 85(9), 18–19.
- Lovallo, D., & Sibony, O. (2010, March). The case for behavioral strategy. *McKinsey Quarterly*.
- Spetzler, C., Winter, H., & Meyer, J. (2016). *Decision quality: Value creation from better business decisions*. John Wiley & Sons.
- Tetlock, P. E., & Gardner, D. (2015). *Superforecasting: The art and science of prediction*. Crown Publishers.

Appendix A: The Six-Element Check

Before committing to any tier two or tier three decision, answer each question in writing. A weak answer to any one of them identifies the weakest link, and the weakest link sets the quality of the whole decision.

- **Frame.** Can I state, in one sentence, what is being decided, and would a thoughtful outsider agree it is the right question at the right scope?
- **Alternatives.** Are there at least three genuinely different options on the table, and what would I do if the current favourite were unavailable?
- **Information.** What is the base rate for decisions of this kind, which of my key facts would survive hostile cross-examination, and what do I wish I knew?
- **Values.** Were the criteria and their weights written down before evaluation began, and do they reflect what actually matters rather than what is easy to measure?
- **Reasoning.** Is the argument written down in a form another person could check, and does the conclusion survive reasonable changes to the weights and probabilities?
- **Commitment.** Do the people who must carry this out believe in it, and is there a named owner, a first step, and a date?

Appendix B: The Premortem in Five Steps

- Assemble everyone with knowledge of the decision, shortly before final commitment, for one hour.
- Announce the premise: it is two years from now, the decision was taken, and it has failed completely. This is not a possibility to debate but a fact to explain.
- Each participant independently writes, for five to ten minutes, the history of the failure: what went wrong, in what order, and why.
- Go around the room, collecting one reason per person per round, until every reason is on the table. No rebuttals during collection.
- Sort the reasons into three groups: those that change the decision, those that change the plan, and those that define early warning signals to monitor. Amend accordingly.

Appendix C: A Worked Example

The example is fictional but built from familiar parts, and this time the numbers are shown, because a framework becomes real at the moment arithmetic enters the room. A family office holds a thirty-one percent stake in a renewable natural gas venture it backed seven years ago with an investment of eleven million dollars, when the company was four employees and a permit application. The venture now operates two facilities converting agricultural waste into pipeline-grade gas, with a third under construction. On a Tuesday in March, an email arrives from the chief executive of a large utility. The utility proposes to buy the company outright at an enterprise value of two hundred and ten million dollars, which after debt puts the family's stake at roughly forty-six million, a little over four times invested capital. The offer expires in forty-five days.

The first meeting is euphoric, which is precisely the danger. The patriarch's instinct is to accept before the buyer changes its mind. His daughter, who sits on the venture's board, wants to refuse outright: the third facility will nearly double capacity, and selling now means selling that growth to someone else. The office's investment director sits between them with a spreadsheet that could justify either position. In an ordinary process, whoever argues best over the next fortnight wins, and the family calls the result a decision. Instead, the office declares it a tier three matter and applies the framework.

Framing comes first, and it changes the question. The decision is not 'accept or decline forty-five days from now.' Written out, it becomes: 'What role should this holding play in the family's next decade, and is this offer the best available way to advance it?' The wider frame surfaces facts the celebration had skipped. At forty-six million dollars against total family wealth of about two hundred and forty million, the position has grown to nineteen percent of everything the family owns, nearly twice the ten percent concentration ceiling the office wrote into its own policy three years earlier. The patriarch is sixty-seven and has begun to speak about simplifying. The daughter's board seat is her main operating role and part of how she is preparing to lead the office. This is not only a valuation question. It never was.

The reframed question generates five genuine alternatives where there had been two. Accept in full. Decline and continue. Sell roughly half the stake to the buyer while negotiating governance protections for the remainder. Decline this offer but run a structured process to test the market, using the bid as a floor. Defer, tying any decision to the third facility reaching commercial operation, scheduled in fourteen months. The vanishing-options test earns its keep here: asked what they would do if selling were simply impossible, both the patriarch and the investment director admit they would be reducing concentration some other way, which tells everyone the true underlying decision is about exposure, not about this buyer.

Information gathering starts from the outside view, and the outside view stings. The office assembles what it can document on comparable ventures that declined credible strategic offers: trade press, broker studies, the fund data its advisors can access. Of the forty-one cases it can reconstruct, fourteen

achieved a better exit within five years. Among the remaining twenty-seven, the median eventual outcome came in roughly thirty percent below the bid that had been declined. The venture's own projections are then restated as ranges rather than points, and the chief executive is asked, in writing, for the ten percent downside case. His answer is worth the exercise by itself: a plausible fifteen-month slip in the interconnection permit for the third facility, to which he assigns about one chance in four, and which would cut the facility's contribution to enterprise value by something like thirty million dollars. That risk had never appeared in a board deck. Nobody concealed it; it had simply never been anyone's job to write the pessimistic history in advance.

Values are set before any scoring, and setting them forces the family's first honest argument, which is the point. After debate, six criteria are written down and weighted out of one hundred: after-tax proceeds at twenty-five, reduction of concentration risk at twenty-five, alignment with the office's energy thesis at fifteen, the daughter's operating role at fifteen, obligations to the founding management team at ten, and the patriarch's stated wish for simplicity at ten. The weighting session produces two discoveries. The patriarch finds that simplicity matters more to him than the last increment of price, and says so aloud for the first time. The daughter finds, somewhat to her own surprise, that what she wants to protect is her role in energy investing, not this company specifically.

Each of the four principals then scores the five alternatives against the six criteria independently, in writing, on a scale of one to ten, before any discussion. The averaged results, weighted, come out as follows.

P is after-tax proceeds (weight twenty-five), C is concentration reduction (twenty-five), T is thesis alignment (fifteen), R is the daughter's role (fifteen), M is management obligations (ten), and S is simplicity (ten). Scores are group averages on a ten-point scale; totals are weighted out of ten. The table does not make the decision, and no one pretends it does. What it makes is the conversation. The race is between the partial sale at 7.10 and the full sale at 6.80, three tenths of a point apart, and a five-minute sensitivity check shows exactly what the gap depends on: move ten points of weight from the daughter's role to the patriarch's simplicity and the full sale overtakes at 7.50 against 7.20. In plain language, the choice turns on how the family weighs one generation's wish to step back against the other's need to step up, which is a question they can now discuss directly instead of smuggling it inside arguments about valuation.

The divergences in the individual score sheets matter more than the averages. The daughter has scored the partial sale nearly two points below everyone else, and pressed for a reason, she produces one nobody had raised aloud: as a minority holder without a path to control, the family would keep the risk of the buyer's decisions while losing the power to affect them. That single observation reshapes the negotiating instructions for the partial-sale option, converting hopes into conditions: a board seat, a list of reserved matters, and a put right exercisable after year three at a defined valuation formula rather than a vague promise of fairness.

A premortem follows on the leading option. It is two years on and the protected partial sale has failed: each participant writes the story. Three failure histories recur. The buyer, post-acquisition, starves the third facility to feed its own capital plan. The put right proves unenforceable in practice because the valuation mechanism was left to future agreement. The founding chief executive, demoralised at becoming a division manager, leaves within a year, and the venture's edge walks out with him. Each of the three becomes either a contract term or a named early-warning signal with a monitoring owner: a committed capital schedule for the facility written into the shareholders' agreement, the valuation formula fixed at signing, and a retention package for the chief executive negotiated as part of the deal. None of this had appeared in any prior meeting, because prior meetings had been about whether to sell, not about how selling fails.

The decision, taken in week five, is the partial sale on protected terms: half the stake sold for about twenty-three million dollars, concentration brought back inside the ten percent ceiling, and a defined mandate, funded at five million, for the daughter to build the office's next energy position. The one-page log records the frame, the five alternatives, the base-rate evidence of fourteen better exits in forty-one cases, the weights, the scoring table, the sensitivity result, the premortem findings, and two revisit triggers. Fourteen months later the third facility is four months late on interconnection, exactly the kind of slip the base rate and the downside case had warned about, and the family reads the log with something close to gratitude. Whether the retained stake ultimately prospers, no one yet knows. What the family does know is that nothing about the choice was left to charisma, momentum, or the loudest voice on a Tuesday in March, and that when they review it, they will be able to judge the process on its merits, whatever luck decides to do with the outcome.