farmBI

Good Decisions

Ten cross-industry drivers of repeatable decision quality

Document type	Research brief
Version	1.0
Date	May 2026
Status	Published
Prepared by	farmBI Limited

Purpose

To identify the core human and organisational drivers that allow operators, leaders, teams and enterprises to make better decisions repeatedly under uncertainty. The emphasis is deliberately pan-industry rather than farm-specific, drawing on decision science, behavioural economics, naturalistic decision-making, high-reliability organisations, evidence-based management, forecasting research and team learning literature.

> **Core premise**
>
> Good decision-makers are not always right; they use processes that improve expected outcomes over repeated decisions.
> Decision quality should be judged separately from outcome quality, because real-world outcomes are noisy.
> Sustained high performance usually reflects a decision operating system: objectives, signals, mental models, social challenge and learning loops working together.

Executive summary

This brief develops ten core drivers of good decision-making across industries. The central argument is that repeatable decision quality is not a single trait such as intelligence, experience or confidence. It is a system of habits and conditions that help decision-makers define what they are optimising for, understand the real situation, reason under uncertainty, challenge their assumptions and learn from outcomes.

The most useful distinction is between decision quality and outcome quality. A decision can be sound and still produce a poor result when weather, markets, biology, geopolitical events, customer behaviour or competitor actions move against the decision-maker. Conversely, a poor decision can sometimes be rewarded by luck. Over time, however, processes that improve signal detection, calibration, challenge and feedback should produce better expected results.

The literature reviewed here suggests five broad clusters. First, good decision-makers establish clear objectives and trade-offs. Second, they improve input quality: situation awareness, evidence discipline and causal models. Third, they improve thinking quality: probabilistic reasoning, cognitive reflection and the careful use of intuition. Fourth, they improve social quality: dissent, psychological safety and deference to expertise. Fifth, they improve learning quality: premortems, staged options, decision logs and debriefs.

For complex operating environments, the strongest performers are likely to be those who combine local knowledge with structured decision hygiene. They are neither purely data-driven nor purely intuitive. They use data to test judgment, judgment to interpret data, and feedback loops to improve both.

The ten drivers at a glance

Driver	Core question	Primary disciplines	Failure mode when weak
1. Clear objective function	What are we actually optimising for?	Decision analysis, strategy	Activity without strategic coherence
2. Situation awareness	What is really happening now?	Human factors, HROs	Decisions based on stale or partial reality
3. Causal mental models	Why is it happening and what changes what?	Systems thinking, causal reasoning	Mistaking correlation, symptoms or anecdotes for causes
4. Evidence discipline	What evidence should inform the call?	Evidence-based management	Overweighting opinion, precedent or the loudest voice
5. Probabilistic thinking	How likely are the plausible futures?	Forecasting, statistics	False certainty and poor calibration
6. Cognitive reflection	What might my first answer be missing?	Behavioural economics	Anchoring, confirmation bias, sunk cost and overconfidence
7. Earned intuition	Is this a domain where gut feel is trustworthy?	Naturalistic decision-making	Using intuition outside its valid feedback environment
8. Constructive dissent	Who can safely challenge the view?	Team learning, psychology	Groupthink, silence and late surfacing of risk
9. Premortem and option design	How could this fail, and how do we preserve options?	Risk, strategy, complexity	Big irreversible bets without testing assumptions
10. Learning loops	How do we turn outcomes into better future judgment?	AARs, debriefs, learning science	Repeating errors and confusing luck with skill

Framing: decision quality versus outcome quality

A decision should be evaluated by the quality of the process and information available at the time, not only by the eventual result. This distinction is central in noisy domains. In business, sport, medicine, investing, military command, aviation, logistics and agriculture, a decision-maker often acts before all uncertainty has resolved. If the only measure is outcome, luck can disguise poor reasoning and misfortune can punish good reasoning.

Decision Quality, as described by Strategic Decisions Group and related decision-science educators, breaks high-quality decisions into requirements such as a useful frame, alternatives, meaningful information, clear values and trade-offs, sound reasoning and commitment to action

(Strategic Decisions Group, n.d.; Decision Education Foundation, n.d.). That framework is useful because it treats the decision as a designed process rather than a moment of inspiration.

This brief extends that idea with evidence from human factors, behavioural economics, forecasting, naturalistic decision-making and organisational learning. The ten drivers are not mutually exclusive. They reinforce one another. For example, a clear objective function improves evidence selection; evidence discipline improves causal models; probabilistic thinking improves premortems; and debriefs improve future intuition.

> **Working definition used in this brief**
>
> A good decision is a choice made through a process that is more likely than the available alternatives to produce the intended result over repeated decisions, given the uncertainty and information available at the time.
> A good decision process reduces avoidable error; it cannot remove uncertainty.

1. Clear objective function

Good decision-makers start by clarifying what “good” means. Without a clear objective function, people optimise for whatever is most visible, urgent, emotionally salient or politically safe.

In decision analysis, framing is not a cosmetic first step; it determines what information matters, which options are considered and how trade-offs are handled. Strategic Decisions Group describes decision quality as including the right frame, meaningful alternatives, useful information, clear values and trade-offs, sound reasoning and commitment to action (Strategic Decisions Group, n.d.). The practical implication is that a decision should not start with “What should we do?” but with “What are we trying to achieve, over what time horizon, with what constraints and unacceptable risks?”

A weak objective function produces hidden substitution. A business says it wants long-term value but rewards quarterly volume. A hospital says it wants safety but punishes error reporting. A trader says they want disciplined returns but behaves as if the objective is to avoid the pain of a losing day. Across industries, people drift into optimising for the metric that is most immediate, easiest to measure or most socially rewarded.

The better pattern is explicit trade-off design. Strong decision-makers name the primary objective, the secondary objectives and the constraints. They distinguish “must not breach” limits from “nice to improve” measures. They also revisit the objective when context changes. This matters because complex systems often contain competing goals: efficiency versus resilience, growth versus optionality, yield versus risk, speed versus quality, local optimisation versus whole-system performance.

Observable behaviours

- Writes the decision objective before analysing options.
- Separates goals, constraints, preferences and non-negotiables.
- Defines the time horizon: immediate result, season or year result, and long-run enterprise effect.
- Makes trade-offs visible instead of pretending they do not exist.

Decision risk

The decision process becomes reactive. People collect data without knowing what would change the decision, compare options against inconsistent criteria, or confuse a successful tactic with a successful strategy.

2. Situation awareness

Good decision-makers maintain an accurate, current view of the operating environment. They know what is happening, what it means and what is likely to happen next.

Situation awareness is heavily studied in aviation, defence, emergency response, healthcare and human factors. Endsley's influential model defines situation awareness as three levels: perception of relevant elements, comprehension of their meaning and projection of their future status (Endsley, 1995). This matters because many errors begin before analysis: the decision-maker never saw the right cue, saw it too late, or failed to connect it to the wider pattern.

High-reliability organisation research makes a related point. Organisations operating in high-risk settings seek "collective mindfulness" through sensitivity to operations, reluctance to simplify, preoccupation with failure, deference to expertise and commitment to resilience (Veazie et al., 2019). These principles are useful beyond safety-critical industries because they describe how strong operators avoid complacency: they treat small anomalies, near misses and weak signals as data rather than as noise to be dismissed.

Good situation awareness is not just having more dashboards. It requires choosing which signals matter, updating frequently enough, and keeping frontline or local observations connected to decision-makers. It also requires understanding lag. Some indicators describe what already happened, while others give a clue about future constraints. Strong decision-makers do not wait until the result is obvious to everyone; they detect changes early enough to preserve useful options.

Observable behaviours

- Maintains a live picture of key operating constraints and leading indicators.
- Distinguishes signal from noise but does not dismiss anomalies too quickly.
- Uses frontline observations and expert local knowledge, not just summary reports.
- Asks: "What would we expect to see if our current view were wrong?"

Decision risk

Decisions are made from a stale mental map. The decision-maker may be technically competent but responding to last month's system rather than today's system.

3. Causal mental models

Good decision-makers understand the causal structure of the system well enough to anticipate second-order effects, delays and unintended consequences.

Causal reasoning is central to decision-making because decisions are interventions: the decision-maker is not merely describing a system but trying to change it. Cognitive science

research on mental models describes how people represent causal assertions and use those representations for deductive, inductive and abductive reasoning (Khemlani, Barbey & Johnson-Laird, 2014). In plain terms, strong decision-makers have better “if this, then what?” models.

Systems thinking adds that real-world systems include feedback loops, delays, stocks, flows and leverage points. Meadows argued that leverage points are places in a system where a small change can produce large effects, but also warned that people often identify them incorrectly (Meadows, 1999). That warning is valuable: operators often see visible symptoms and act on them, when the real leverage sits elsewhere in the system.

The Cynefin framework also helps explain why causal models must match the context. Snowden and Boone distinguish clear, complicated, complex and chaotic domains; in complex domains, cause and effect may only be clear in hindsight, so leaders need experiments that are safe to fail rather than overconfident prediction (Snowden & Boone, 2007). The lesson is not “never model”; it is “model humbly, and choose the right type of action for the kind of system you are in.”

Observable behaviours

- Maps the drivers of performance, constraints, feedback loops and lagged effects.
- Distinguishes correlation from causation and symptoms from root causes.
- Tests whether a chosen intervention addresses the bottleneck or only the visible problem.
- Recognises when a system is complex enough to require experimentation rather than prediction alone.

> **Decision risk**
> The decision-maker becomes a collector of indicators rather than an interpreter of systems. Data may increase confidence without improving understanding.

4. Evidence discipline

Good decision-makers deliberately combine multiple sources of evidence and appraise the quality of each source before acting.

Evidence-based management defines good practice as making decisions through the conscientious, explicit and judicious use of the best available evidence from multiple sources (Barends, Rousseau & Briner, 2014). Those sources typically include scientific evidence, organisational data, professional expertise and stakeholder values or concerns. The value of this model is that it does not worship academic research, internal data or experience alone. It asks how each source should be weighted for the decision at hand.

The CIPD’s evidence-based practice guidance makes the same practical point: evidence is information that supports or contradicts a claim, assumption or hypothesis, and it should come from multiple sources (CIPD, 2025). This is especially relevant in messy operating environments where data may be incomplete, expert opinion may be biased, and published research may not perfectly fit the local context.

Evidence discipline also means deciding in advance what evidence would change the decision. This prevents analysis from becoming theatre. Weak decision processes often gather information to justify a preferred option; strong processes gather information to discriminate

between plausible options. They actively look for disconfirming evidence, missing base rates and external reference classes.

Observable behaviours

- Uses more than one evidence source: data, experience, research and stakeholder insight.
- Checks data provenance, freshness, measurement error and relevance.
- Looks for disconfirming evidence, not just confirming evidence.
- States what evidence would change the decision before further analysis begins.

> **Decision risk**
>
> Decisions become over-dependent on anecdotes, benchmarking, habit, authority, dashboards or a single impressive data point.

5. Probabilistic thinking

Good decision-makers think in probabilities, ranges and scenarios rather than binary predictions. They are calibrated rather than merely confident.

Forecasting research shows that judgment can improve when people express beliefs probabilistically, receive feedback, collaborate constructively and track accuracy. In geopolitical forecasting tournaments, Mellers and colleagues found that probability training, team collaboration and tracking improved both calibration and resolution (Mellers et al., 2014). Forecasting tournaments also create a discipline of making beliefs explicit enough to be scored.

Probabilistic thinking is powerful because many important decisions are about expected value, risk exposure and timing rather than certainty. A decision-maker who says “this will happen” has given little information about uncertainty. A decision-maker who says “I think this has a 65% probability, here is the base rate, and here are the signals that would move me to 80% or 40%” has created a decision that can be challenged and learned from.

This does not require false precision. In many real-world settings, rough calibrated ranges are more useful than point forecasts. The habit matters: separate the base case from upside and downside; estimate likelihood and consequence separately; and track whether predictions are consistently too optimistic, too pessimistic or too narrow.

Observable behaviours

- Uses likelihood ranges, not just labels such as “likely” or “risky”.
- Separates probability from impact.
- Uses base rates and outside views before adjusting for local specifics.
- Tracks forecast accuracy over time to improve calibration.

> **Decision risk**
>
> The decision-maker oscillates between overconfidence and indecision. Uncertainty is either ignored or treated as a reason not to act.

6. Cognitive reflection and bias control

Good decision-makers can pause, inspect their first answer and apply decision hygiene to reduce predictable cognitive errors.

Frederick's Cognitive Reflection Test was designed to measure the ability or disposition to resist the first answer that comes to mind when that answer is intuitively appealing but wrong (Frederick, 2005). The broader lesson is not that slow thinking is always better; it is that high-stakes or unfamiliar decisions often require a deliberate interruption of automatic judgment.

Kahneman, Lovallo and Sibony argue that executives often understand biases in general but fail to recognise them in their own live decisions. Their checklist approach shifts attention from trying to make individuals unbiased to improving the decision process around recommendations (Kahneman, Lovallo & Sibony, 2011). This is a critical distinction. Bias control is more reliable when embedded in routines, checklists, independent estimates and challenge processes than when it relies on personal self-awareness alone.

Common errors include anchoring on the first number, searching for confirming evidence, being overconfident in narrow forecasts, escalating commitment because of sunk cost, overweighting recent events, and preferring action that protects ego or status. Strong decision-makers do not assume they are immune. They use process design to make errors easier to detect.

Observable behaviours

- Identifies the anchor, the outside view and the main disconfirming evidence.
- Asks whether the decision would look different if the sunk cost were zero.
- Uses independent estimates before group discussion where possible.
- Builds checklists for repeated high-consequence decisions.

> **Decision risk**
> Confidence becomes a poor proxy for correctness. The decision-maker may be persuasive, experienced and wrong for predictable reasons.

7. Earned intuition

Good decision-makers know when intuition is a source of expertise and when it is merely a feeling of fluency.

Kahneman and Klein reached a useful synthesis on professional intuition. They argued that skilled intuition is more likely to be valid when the environment has sufficiently regular patterns and the decision-maker has had adequate opportunity for feedback and learning (Kahneman & Klein, 2009). This explains why intuition can be excellent in some domains and dangerous in others.

Naturalistic decision-making research, including Klein's work on recognition-primed decision-making, studies how experienced people make effective decisions under time pressure, high stakes, ambiguity and changing conditions (Klein, 1998; Klein, 2008). Experts often do not compare a long list of options. They recognise a situation as typical or atypical, simulate a plausible action mentally, and adjust if the action will not work.

The practical challenge is to distinguish earned intuition from overfitted intuition. Intuition is more trustworthy when feedback is frequent, clear and causally connected to the decision. It is less

trustworthy when outcomes are delayed, noisy, politicised or heavily affected by luck. A strong decision-maker therefore asks: “Is this a domain where experience has trained me, or a domain where experience may be misleading me?”

Observable behaviours

- Trusts intuition more in high-validity, high-feedback environments.
- Uses intuition as an input, then tests it against evidence and alternatives.
- Notices when a situation is novel enough that pattern recognition may fail.
- Develops intuition deliberately through repeated decisions, feedback and reflection.

> **Decision risk**
> People either reject valuable tacit expertise or over-trust gut feel in environments where feedback has not trained good judgment.

8. Constructive dissent and psychological safety

Good decision-makers create conditions where knowledgeable challenge, bad news and minority views can surface early enough to matter.

Edmondson defines psychological safety as a shared belief that a team is safe for interpersonal risk-taking, and her work links psychological safety with learning behaviour in teams (Edmondson, 1999). This does not mean comfort, consensus or politeness. It means people can ask questions, admit uncertainty, raise concerns and challenge assumptions without fear of humiliation or punishment.

High-reliability organisation principles also emphasise deference to expertise: in critical situations, insight should flow from the person with the most relevant knowledge, not merely the person with the highest status (Veazie et al., 2019). This is relevant to any complex enterprise. The people closest to the work often see weak signals before senior decision-makers do.

Constructive dissent improves decision quality by increasing the diversity of information and reducing social distortion. However, dissent must be structured. Unstructured debate can become personality conflict or political theatre. Better routines include red teams, devil’s advocates, independent pre-reads, premortems, “what would make this wrong?” questions and explicit invitation of the most informed opposing view.

Observable behaviours

- Invites challenge before commitment, not after failure.
- Separates critique of the idea from critique of the person.
- Asks lower-status or frontline experts what leaders may be missing.
- Rewards early surfacing of risk rather than punishing bad news.

> **Decision risk**
> Teams slide into groupthink, silence, false consensus or delayed escalation. The decision looks aligned because dissent has gone underground.

9. Premortem and option design

Good decision-makers imagine failure before commitment and design decisions to preserve learning, reversibility and options where possible.

Gary Klein's premortem method asks a team to imagine that a plan has failed and then work backward to identify plausible reasons for the failure (Klein, 2007). The technique is valuable because it legitimises dissent before the decision becomes personally or politically owned. It also changes the emotional tone: instead of saying "I disagree with the plan," participants can say "If this failed, here are the ways it might have happened."

Premortems connect naturally to option design. In uncertain environments, the best decision is often not the largest commitment to the most attractive forecast. It may be a staged decision that buys information, limits downside, keeps alternatives open and creates trigger points for escalation, pause or exit. Decision quality frameworks similarly emphasise feasible alternatives and clear trade-offs, not just analysis of a favoured option (Strategic Decisions Group, n.d.).

Complexity research reinforces this logic. In complex domains, Snowden and Boone argue for probing and safe-to-fail experiments rather than assuming that analysis can fully specify the right answer in advance (Snowden & Boone, 2007). Strong decision-makers therefore ask not only "What is the best option?" but "What option gives us the best learning-adjusted path under uncertainty?"

Observable behaviours

- Runs premortems before major commitments.
- Identifies leading indicators that would show the decision is failing.
- Stages investment or operational commitment where uncertainty is high.
- Builds explicit stop, continue, scale or pivot triggers.

> **Decision risk**
>
> Organisations make large irreversible bets too early, then defend the decision because reversing it would be embarrassing.

10. Learning loops

Good decision-makers turn decisions into data for future decisions. They learn faster because they record expectations, review outcomes and update mental models.

Learning loops are the mechanism that converts experience into expertise. Without them, experience can merely reinforce stories. A decision log records what was decided, what was expected, the evidence used, the probability estimates, the assumptions and the review trigger. When the outcome is known, the decision-maker can compare forecast to result and distinguish luck, execution failure, model error and environmental change.

Debriefing and after-action review research supports the value of structured reflection. Tannenbaum and Cerasoli's meta-analysis concluded that properly conducted debriefs can improve individual and team performance by roughly 20% to 25% (Tannenbaum & Cerasoli, 2013). Later meta-analytic work on after-action reviews has also examined the task and training characteristics that influence effectiveness (Keiser & Arthur, 2021).

The most effective learning loops are timely, specific and psychologically safe. They avoid blame-first explanations and hindsight certainty. They ask: What did we intend? What actually happened? Why did it happen? What should we sustain, change or test next time? Over time, this improves calibration, causal models, intuition and team trust.

Observable behaviours

- Maintains a decision log for repeated or high-consequence choices.
- Reviews decisions against prior expectations, not rewritten memories.
- Separates process quality, execution quality and luck.
- Converts lessons into updated checklists, thresholds, models or operating routines.

> **Decision risk**
> People repeat mistakes, attribute success to skill and failure to external events, and lose the chance to compound decision quality over time.

Integrated framework: the decision operating system

The ten drivers are most useful when treated as an integrated operating system rather than a checklist of personal virtues. A decision process begins with framing, objective clarity and constraints. It then builds situation awareness and evidence quality. It uses causal models and probabilistic reasoning to generate expectations. It applies cognitive reflection and earned intuition to interpret what the evidence means. It invites constructive dissent and premortem thinking before committing. Finally, it records expectations and reviews outcomes to improve future decisions.

This operating-system view also explains why sustained high performance is rare. It requires the discipline to do basic things repeatedly when conditions are busy, emotionally charged or ambiguous. The edge compounds because each cycle improves the next: better logs improve calibration; better calibration improves premortems; better premortems improve option design; better option design produces safer learning; safer learning improves intuition.

Maturity indicators

Maturity level	Typical pattern	Decision risk	Upgrade move
Ad hoc	Decisions rely on instinct, urgency and precedent.	Luck is mistaken for skill; errors are hard to diagnose.	Write objectives and assumptions before important decisions.
Informed	Some data and expert input are used, but process varies.	Evidence may justify a preferred answer.	Define what evidence would change the decision.
Disciplined	Options, probabilities, risks and trade-offs are explicit.	Process can become bureaucratic if not tied to action.	Use premortems and staged triggers.
Learning system	Decision logs, debriefs and updated models are routine.	Requires cultural safety and time discipline.	Track calibration and embed lessons into future routines.

Practical checklist for evaluating a major decision

Checklist area	Question
Objective	What outcome are we optimising for, over what time horizon? What are the hard constraints?
Frame	Is this the right decision to be making, or have we framed it too narrowly?
Situation	What is happening now? Which indicators are leading, lagging or noisy?
Causality	What causal model explains the situation? What bottleneck or leverage point are we acting on?

Checklist area	Question
Evidence	What evidence supports the decision? What evidence contradicts it? How reliable is each source?
Alternatives	Have we compared genuinely different options, including doing less, staging or waiting?
Probabilities	What are the base case, downside and upside scenarios? What probabilities do we assign?
Bias check	What are we anchored to? What sunk costs, incentives or ego risks may distort the call?
Dissent	Who disagrees, who is silent and who has the most relevant expertise?
Premortem	Imagine the decision failed. What are the three most plausible reasons?
Triggers	What leading indicators would cause us to stop, continue, scale or pivot?
Learning	How will we record this decision and review whether our assumptions were right?

Conclusion

The cross-industry evidence points to a consistent theme: good decisions are engineered. They come from repeatable processes that improve framing, signal quality, reasoning, challenge and learning. The best operators are not purely intuitive and not purely analytical. They integrate experience, data, judgement and disciplined reflection.

For any enterprise seeking sustained top-quintile performance, the goal should be to build a decision system that compounds. Each decision should leave behind a clearer model, better calibration and sharper intuition for the next decision. That is the mechanism by which good decisions become repeatable performance rather than isolated wins.

References

Barends, E., Rousseau, D. M., & Briner, R. B. (2014). Evidence-Based Management: The Basic Principles. Center for Evidence-Based Management. cebma.org

CIPD. (2025). Evidence-based practice for effective decision-making. cipd.org

Decision Education Foundation. (n.d.). What is Decision Science? decisioneducation.org

Edmondson, A. C. (1999). Psychological Safety and Learning Behavior in Work Teams. Administrative Science Quarterly, 44(2), 350–383.

Edmondson, A. C. (2002). Managing the risk of learning: Psychological safety in work teams. Harvard Business School Working Paper.

Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. Human Factors, 37(1), 32–64.

Frederick, S. (2005). Cognitive Reflection and Decision Making. Journal of Economic Perspectives, 19(4), 25–42.

Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. American Psychologist, 64(6), 515–526.

Kahneman, D., Lovallo, D., & Sibony, O. (2011). Before you make that big decision. Harvard Business Review, 89(6), 50–60.

Keiser, N. L., & Arthur, W., Jr. (2021). A meta-analysis of the effectiveness of the after-action review (or debrief) and factors that influence its effectiveness. Journal of Applied Psychology, 106(7), 1007–1032.

Khemlani, S. S., Barbey, A. K., & Johnson-Laird, P. N. (2014). Causal reasoning with mental models. Frontiers in Human Neuroscience, 8, 849.

Klein, G. (1998). Sources of Power: How People Make Decisions. MIT Press.

Klein, G. (2007). Performing a Project Premortem. Harvard Business Review.

Klein, G. (2008). Naturalistic Decision Making. Human Factors, 50(3), 456–460.

Meadows, D. H. (1999). Leverage Points: Places to Intervene in a System. The Sustainability Institute.

Mellers, B., Ungar, L., Baron, J., et al. (2014). Psychological strategies for winning a geopolitical forecasting tournament. Psychological Science, 25(5), 1106–1115.

Snowden, D. J., & Boone, M. E. (2007). A Leader's Framework for Decision Making. Harvard Business Review, 85(11), 68–76.

Strategic Decisions Group. (n.d.). Decision Quality. sdg.com

Tannenbaum, S. I., & Cerasoli, C. P. (2013). Do team and individual debriefs enhance performance? A meta-analysis. Human Factors, 55(1), 231–245.

Veazie, S., Peterson, K., & Bourne, D. (2019). Evidence Brief: Implementation of High Reliability Organization Principles. U.S. Department of Veterans Affairs.