

RED BEAR ARTIFICIAL INTELLIGENCE

ABU DHABI, UNITED ARAB EMIRATES

Building a Decision-Making Culture

Judge the decision, not the result: decision governance as the evidence a fair culture needs

People decide well when they believe they will be judged on their reasoning rather than on their luck, and that belief cannot survive in an organisation that keeps no record of what anyone was thinking at the time. This paper argues that a decision-making culture is a matter of evidence rather than intent, sets out five practices that need no software, and describes the record Red Bear: Enterprise Decision Intelligence keeps so that a review a year later can be honest.

Empower. Trust. Verify.

White paper EDI-WP-05 · Red Bear white paper series · Version 2.0

26 September 2026

Summary

Most organisations do not have a decision-making problem. They have a problem with what happens after a decision turns out badly. Everything else, the escalation that did not need to happen, the proposal nobody put forward, the manager who waits for cover before committing to something well inside her authority, follows from that, and no amount of delegated authority on paper will change it while the answer to "what happens if this goes wrong" remains unclear.

A decision is a judgement made under uncertainty, with the information available at the time. An outcome is what the world subsequently did. In any business with real uncertainty in it they come apart regularly: a carefully reasoned decision produces a poor result, and a careless one gets lucky. Annie Duke calls judging the decision by its result *resulting*, and the word is useful because executives recognise the behaviour in themselves the moment they hear it named. An organisation that judges by outcome alone does not teach its people to decide well. It teaches them to avoid visible risk, which is a different skill, easier to learn, and corrosive in proportion to how well it is learned.

The clearest instrument for this is the escalation rate. When people consistently push decisions upward that they are entitled to make, they are not confused about their authority. They are telling you what they expect to happen if it goes wrong. Escalation is usually treated as a workflow problem; it is more useful read as a measurement of felt safety, and it is one of the few measurements of culture that arrives without a survey.

Here is the hard part. Telling people you will judge the decision rather than the outcome is worthless unless you can reconstruct what they knew and thought at the time. Without a contemporaneous record, every review is hindsight wearing a fair-minded expression, and everybody in the room knows it. Fair review is not a matter of intent. It is a matter of evidence, and the evidence is not currently kept.

The operating model this paper argues for is small: separate the judgement from the result and review it by class rather than by incident; take independent assessments before discussion; give dissent a legitimate place through the pre-mortem; classify the problem as generic or exceptional before deciding who owns it; and weight input by demonstrated track record, carefully. None of it requires software. What does require something is the part where you go back, a year later, and review honestly. Red Bear: Enterprise Decision Intelligence (formerly GuardrAll) is described here as the implementation of that record, and the paper concedes most of the objection that software cannot fix a culture.

WHO SHOULD READ THIS

- Chief executives and chief people officers
- Chief risk officers and heads of internal audit
- Leaders who have announced a no-blame culture and wondered why nothing changed
- Anyone whose decisions will be recorded

WHAT YOU WILL BE ABLE TO DO AFTERWARDS

- Read the four cells of the reasoning-and-result table, and say what a resulting culture concludes in each
- Use escalation rate as a measurement of felt safety
- Run the five practices, and know which two carry most of the weight
- State the three observable differences between a record built for learning and one built for blame

1 What it is

Consider a supplier approved at a modest price advantage, after checking the financials, the references and the single-source risk, with the reasoning written down. Eight months later the supplier fails. Was that a bad decision? On the information available, no. The organisation bought an exposure it understood at a price it judged fair, and the exposure came in. If you would take that bet a hundred times because it pays on ninety of them, the nine times it does not are the cost of the ninety. Now consider the same approval made because the incumbent was irritating and the presentation was good, with no reference check and nothing written down. Eight months later the supplier performs well. A resulting culture records this as a success, promotes the method, and quietly increases the amount of money committed this way.

The reasoning	The result	The honest conclusion	What a resulting culture concludes
Sound	Good	*Deserved success.* The process worked and the world cooperated. Do it again	Correct, by accident. The method is rewarded for the wrong reason
Sound	Bad	*Bad luck.* The process worked. Find the assumption that failed and price it better next time	A bad decision. Somebody should have known. Authority is quietly withdrawn
Careless	Good	*Dumb luck.* A near miss. The process is unsafe and got away with it once; the most dangerous cell	A good decision. The method is praised, copied and scaled
Careless	Bad	*Deserved failure.* The process failed. Fix the process	A bad decision. Correct, and usually the only case it gets right for the right reason

The Red Bear framework scores decision quality and outcome quality separately and holds the pair as a queryable quadrant with no combining axis; dumb luck is marked the most dangerous of the four (Oliver, *Empower, Trust, Verify*, Book 1, chapter 2). The asymmetry that makes resulting so persistent is not a failure of character. It is a failure of evidence. Outcomes are visible, dated and attributable: the supplier failed, the write-off has a number. Reasoning is invisible, undated and mostly unrecorded. When the only surviving artefact of a decision is its result, review has nothing else to work with, so it works with the result and calls that rigour.

Drucker supplies the second half of the distinction. His first test is to classify the problem: is it generic, one of a recurring class that should be answered once by a rule, or genuinely exceptional? Most of what organisations treat as a series of unique judgement calls is a generic class being decided repeatedly from first principles, which is both slow and unreviewable. One outcome tells you almost nothing about a judgement made under uncertainty; fifty of the same class tell you a great deal (Drucker, *The Effective Decision*, Harvard Business Review, 1967). The classification is the precondition for learning anything at all, which is why decision class is the most load-bearing attribute in the Red Bear record.

Key points

- A bad decision and a bad outcome are different things; the quality quadrant keeps them apart
- Escalation rate is a measurement of felt safety, available without a survey
- Fair review is a matter of evidence recorded before the outcome was known

2 Why it matters

The first cost is escalation. When a manager with a clear mandate routes a decision upward anyway, two things are consumed: the time of a senior person who adds no information, and the elapsed time of whatever was waiting. Both are measurable today: count the decisions in one class that went above the level authorised to settle them, and note how long each waited. Bain's four dimensions of decision effectiveness, quality, speed, yield and effort, name the one most organisations never look at: effort, because nobody's budget contains a line for it (Blenko, Mankins and Rogers, *Decide & Deliver*, chapter 1).

The second cost is the erosion of judgement itself. Judgement is built by deciding and then finding out, and it compounds. A manager who has settled two hundred pricing exceptions has something a manager who has escalated two hundred does not, and an organisation that discourages the first has not saved risk. It has deferred the acquisition of competence and concentrated what it has in the few people at the top, which is an expensive place to keep it and a fragile one: it retires.

The third cost shows in an exit interview. People who are good at deciding and are not permitted to do it leave, and rarely say that is why. Before they leave they stop proposing things, which is invisible, and they acquire the defensive habits that make a bad outcome survivable: the long email copying everyone, the decision taken jointly by five people so that nobody owns it, the request for a steering group. None of these are irrational. They are correct responses to an accurate reading of the environment.

The limits of measurement are worth stating plainly. Escalation rate, queue time and senior hours consumed can be counted. The decisions that were never proposed cannot. Any paper offering a percentage for the cumulative effect of a workforce that has learned to be careful in the wrong way has invented one. What can be said with confidence is the mechanism: outcome-only review produces risk avoidance, risk avoidance produces escalation, and escalation is countable.

Organisational memory is not documentation. Organisations believe they remember their decisions because they can produce the board paper, the business case, the contract, the minute. Those artefacts record the conclusion, written up after the fact in a system built for the transaction. Actual organisational memory is the other material: that this was tried in 2021 and failed for a reason that still applies; that the alternative supplier was rejected not on price but on a concentration risk that has since changed; that the assumption underneath the forecast was a deliberate conservative choice rather than a belief; that two people argued strongly against, and what they argued. The value is concentrated in the rejected options and the stated assumptions, and those are exactly the things no document retains. The same proposal returns every eighteen months and is argued from the beginning each time. A constraint that was a sensible judgement in one set of conditions hardens into a rule that outlives them. A new executive is told "we tried that" and cannot establish whether it cannot work or was done badly once. This is the fairness problem viewed from a different angle: the same missing asset, held by people who will at some point change jobs.

3 The challenges

Five remedies are commonly attempted, and each fails for a reason that is visible in advance.

Challenge	Why it is hard	What goes wrong without a method
Saying so	Leaders announce that honest mistakes are not blamed, and mean it; staff calibrate on events, not statements	One competent person removed after a defensible decision that went wrong, and every manager who watched has their answer

Challenge	Why it is hard	What goes wrong without a method
The blameless review	It works in engineering because the evidence is complete: logs, timestamps, configuration state; applied to a commercial decision with no contemporaneous record it becomes a reconstruction contest	The most articulate account of what somebody was thinking eight months ago wins: a fluency test, conducted in public
The delegation matrix	It grants permission, and permission is not the binding constraint	The manager who escalates a decision she is authorised to take has read the matrix; she is managing a different risk
Training in psychological safety	It reliably teaches the vocabulary; whether behaviour changes depends on what happens to the next person who takes a risk	The organisation demonstrates that it knows the right words and does not mean them
Improving the escalation process	Faster routing, clearer queues, reminders, a senior service level: all sensible, all reducing the cost of escalating	Treating the symptom efficiently makes it permanent; the work of establishing when escalation is warranted is upstream of any routing improvement
Keeping the record	Reasoning, dissent and the policy of the day must be written before the outcome is known and kept unaltered beside it	Without it, the review a year later cannot be fair, and everyone can see that it cannot

4 How Red Bear: Enterprise Decision Intelligence helps

Five practices, none of which requires a purchase, and all of which are considerably easier to sustain when the reasoning is captured as a matter of routine.

1. **Separate the judgement from the result.** Assess the process, evidence gathered, options considered, assumptions stated, dissent heard, as a thing distinct from what happened, and compare the two across a class rather than arguing about one.
2. **Independent assessments before discussion.** Each participant records a judgement on their own before anyone speaks. Judgements made after discussion are correlated, because people have heard each other; judgements made before it are independent, and their spread is information about how uncertain the question really is (Kahneman, Sibony and Sunstein, *Noise*, part 5).
3. **Give dissent a legitimate place.** The pre-mortem asks the group to imagine the decision has already failed and to say why. Klein reports that prospective hindsight improves people's ability to identify reasons for future outcomes; the figure comes from a small body of laboratory work, so treat it as directional. Its second effect is organisational: it gives the person with a concern a sanctioned way to raise it, and produces a written list of what would have to go wrong (Klein, *Performing a Project Premortem*, Harvard Business Review, 2007).
4. **Classify the problem before assigning it.** Generic problems become policy and are settled at the edge; exceptional ones are escalated deliberately. Most escalation load is generic decisions never recognised as a class.
5. **Weight input by demonstrated track record.** Dalio's believability weighting is a real and defensible idea, and it degrades badly into a caste system if applied without discipline. The guards are specific: track record is per domain, not per person; the basis is visible to the person weighted; it is recomputed from evidence rather than accumulated by reputation; and it never becomes the only route by which a dissenting view reaches a decision, which is one more reason the pre-mortem matters, because it is deliberately unweighted (Dalio, *Principles*, work principles on believability).

The case for the second and third rests on Lovallo and Sibony's study of 1,048 business decisions, which found that process quality mattered roughly six times more than analysis quality in explaining how decisions turned out. The finding should be described accurately: executives reported retrospectively, process and outcome were assessed in the same survey, and it is a strong association rather than a controlled experiment. Even so, the implication is blunt: organisations spend far more on better analysis than on better process, and the returns appear to sit on the other side (Lovallo and Sibony, *The case for behavioral strategy*, McKinsey Quarterly, 2010).

Practice	What Enterprise Decision Intelligence does	Requirements
Reasoning captured at the time	One record per decision across people, agents and applications, with the four Cartesian quadrant answers written by the requester, the policy version in force when the decision was evaluated, and an append-only trail; audit replay reconstructs the record as it stood, not as it reads now	FR-CAP-12, FR-OUT-04, FR-REC-08, FR-REC-09
Independent judgements before discussion	Confidence recorded on a one-to-ten scale per decision and per assumption, immutable once written; noise measured as the spread between independent judgements, not their average	FR-REC-12, FR-OUT-07
Dissent and pre-mortem on the record	The minority view and "what would make this wrong" captured at the time of the decision, attributed and timestamped, for the classes whose standard requires it; a critical decision records the dissent	FR-CAP-11, FR-REC-06
Outcomes attached later	Outcome tracking from a connected system or the accountable owner's review at the class's interval, appended without altering what the record said before the outcome was known	FR-OUT-01, RBP FR-DL-01
Decision quality scored apart from outcome	Process scored without access to any measure series; the quality quadrant queryable with no combining axis; slices by region and decision type, never by name	FR-OUT-05, FR-OUT-08
Classify before assigning	Every candidate classified into routine, discretionary, significant or critical with a confidence; below the threshold, held for a person to classify; the class sets the evidence, memory and Cartesian standards	FR-REC-01, FR-REC-06
Escalation read as a measurement	Bottleneck and latency analysis by class, stage and approver; referral patterns per class, distinguishing repeated referral of one case type from referrals consistently authorised unchanged	FR-RTE-05
Track record, carefully	The earned trust ledger scores trust from each actor's record per decision class and recommends a dial setting to the accountable owner, who decides; it has no path to apply it	FR-DEL-03, FR-DEL-02
Lessons without blame	Patterns found across classes and never attributed to a person; a lesson concerning someone's own decisions is offered as a private coaching conversation, by invitation, with nothing attributable kept	RBP FR-DL-03, RBP FR-DL-04, RBP FR-DC-01, RBP FR-DC-04

Two of those entries carry most of this paper's argument. Dissent and pre-mortem capture put the minority view and the stated failure modes on the record at the moment of the decision, rather than leaving them to be recalled by people who now know how it turned out. Decision quality scoring holds the process separately from the result and compares the two at the level of a class rather than an incident. Keeping the two apart is the entire point. A single blended score that quietly mixes reasoning with luck is a resulting machine with a better interface, and it would do more damage than no score at all.

The earned trust ledger is where believability weighting lives, and the caution applies here with more force than anywhere else in the product, because scoring trust from outcomes is precisely the behaviour this paper argues

against. Hence the score is held per decision class rather than per person, is one input to an owner's judgement rather than a control that moves on its own, and carries process quality alongside outcome. It is a design that can be got wrong, and saying so is more useful than claiming otherwise.

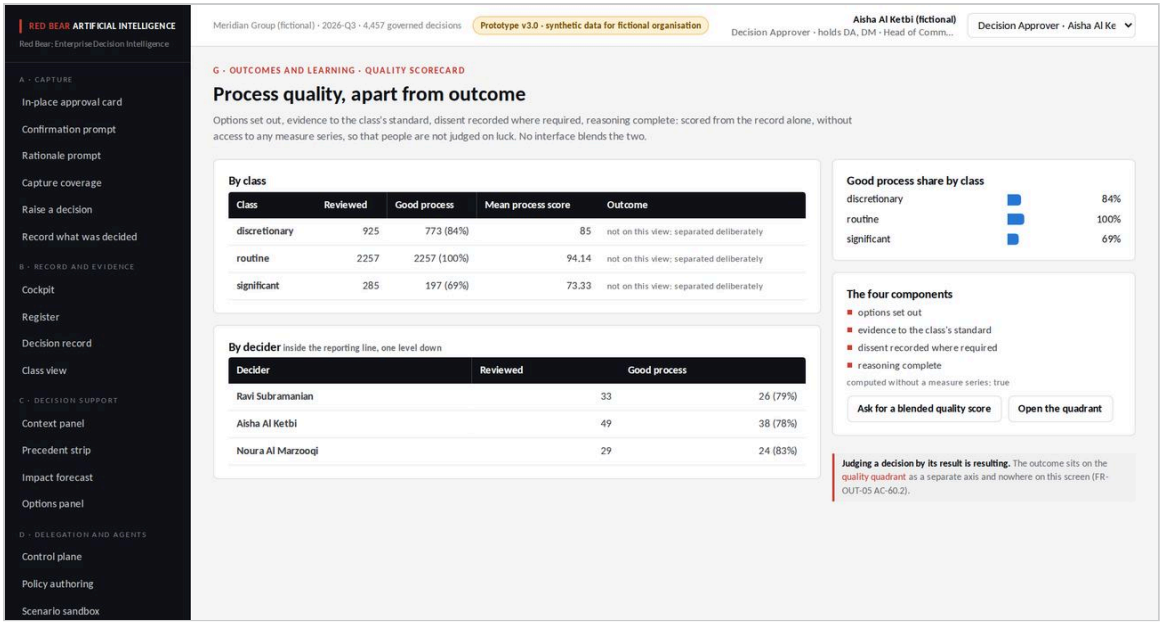
5 How it is used

Step	Who	What happens
1 Count	Sponsor, class owner	Take one decision class with real volume and establish what share went above the level authorised to settle it, and how long each waited: the felt-safety measurement, available now
2 Ask the harder question	Leadership team	What actually happened to the last three people whose defensible decision produced a bad result? If it is not an answer you would want published, nothing else here will take hold until it changes
3 Change what is recorded	Decision Maker	For the next thirty decisions in that class: what is being decided, what is being assumed, and what would make this wrong, written before the outcome is known; the point-of-decision prompt asks this at the moment of commitment
4 Judge independently	Participants	Confidence per decision and per assumption recorded before discussion; notice the spread, which is usually wider than anyone expects
5 Run one pre-mortem	Decision Maker, dissenters	State that the commitment has failed twelve months from now and ask everybody to write down why; the list goes on the record
6 Attach outcomes	Connected system or owner	At the class's interval; the original record is never edited
7 Review in public, and review a success first	Class owner	Pick a decision from the class that worked and review the reasoning rather than the result; do that twice before reviewing a failure

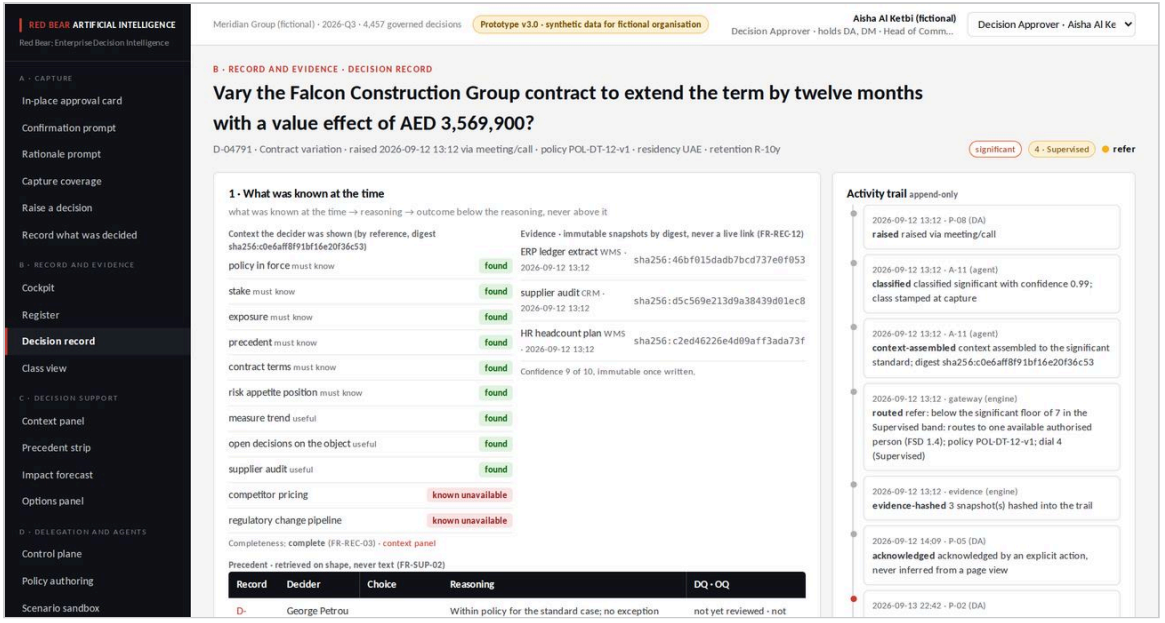
Reviewing a success first is not a presentational trick. It is the only way to demonstrate that the subject of a review is the process, because a review that only ever follows a loss has already told everyone what it is really about.

An illustration. Meridian Group, the invented organisation used in these scenarios, reviews its credit-exception class after twelve months. Of ninety-one decisions, seven produced losses. The review does not begin with the seven. It begins with the whole class, compares the recorded reasoning against the recorded outcomes, and finds that six of the seven followed the same process as the eighty-four that did not lose money. The seventh did not: no alternative was considered and the stated assumption was contradicted by data already in the record. One conversation follows, and it is about process. What the review needed was on the record: the reasoning, recorded before the outcome was known; the dissent, who disagreed and on what grounds; the policy of the day, not today's policy; the class, so fifty decisions can be compared, not one; and the outcome, attached later and kept separate.

6 The user experience



S-41 Quality scorecard: S-41 Quality scorecard. Process quality per class, options set out, evidence to standard, dissent recorded, reasoning complete, computed without access to any outcome series; the outcome shown beside it and never blended. Prototype render on synthetic data for fictional entities.



S-11 Decision record: S-11 Decision record. Question, options, choice and grounds, the four quadrant answers, dissent and pre-mortem, confidence per assumption, the policy version and dial in force, and the outcome appended later. Prototype render on synthetic data for fictional entities.

7 Where the data comes from, and how it is checked

Inputs to the fair review

Input	Source	Check
The decision and its reasoning	Every capture route, terminating in one normalised candidate; the point-of-decision prompt at commitment	Confirm-first: nothing is written until the decider confirms, and an unconfirmed prompt leaves no record (FR-CAP-12)
Independent judgements	Confidence per decision and per assumption, recorded by each participant	Immutable once written; no calibration figure shown below twenty to thirty decisions (FR-REC-12)
Dissent and pre-mortem	Recorded at the time by the people who held them	Attributed and timestamped; required by the class's evidence standard for significant and critical decisions (FR-CAP-11, FR-REC-06)
The policy of the day	Read at evaluation and retained with the record	Replay never re-renders against current settings (FR-REC-09)
The class	Classified with a confidence; customers extend the catalogue	Held for human classification below the threshold (FR-REC-01)
The outcome	A connected system against the decision identifier, or the accountable owner's review	An outcome the product guessed is not evidence (FR-OUT-01)
Escalations and waits	The routing record: stage, approver, target	Attributed to a named stage; referral patterns reported per class (FR-RTE-05)

How it is checked

- **Connectors are read-only and consented.** The decision is captured, not the conversation; consent is recorded per participant, redaction applied before storage, and what is not decision content is discarded once the candidate is formed (RBP FR-CN-01, FR-REC-04).
- **The AI proposes; a person confirms.** A detected decision, a suggested class and a prompted rationale are proposals until a named person confirms them (FR-CAP-08, FR-REC-01).
- **Quality is scored blind to outcomes.** Process quality is computed without access to any measure series, so people are not judged on luck; no interface blends the two (FR-OUT-05).
- **The record is never rewritten.** Outcomes and comparisons are appended; the reasoning stands as written (FR-REC-08).
- **Patterns, never individuals.** Lessons are never attributed to a person; the quality quadrant slices by region and type, never by name (RBP FR-DL-04, FR-OUT-08).
- **Each person sees their own record.** Approver oversight follows the reporting line; coaching is an invitation, not a summons, and nothing attributable is kept from it (FR-PLT-08, RBP FR-DC-01, RBP FR-DC-04).

8 Benefits realised

Every measure is calculated from the organisation's own register against a baseline established before anything was switched on (FR-PLT-05).

Benefit	How it shows up	Measured by
Felt safety rises	People settle decisions they are entitled to settle	Share of decisions in a class escalated above the authorised level, falling
Reviews become fair	Bad luck is distinguished from deserved failure	Share of decisions in each quality-quadrant cell, per class
Judgement develops	Deciders move up the dial on their process record	Dial moves upward per class per year
Dissent is heard	The minority view is on the record before the outcome	Share of significant and critical decisions with dissent or a pre-mortem recorded
Noise becomes visible	Independent judgements spread before discussion	Noise per class: the spread, not the average
Organisational memory persists	Rejected options and stated assumptions survive the people	Precedent retrieved and linked on new decisions in the class

9 Risks managed

The strongest objection to this paper is that culture is made by what leaders do, not by what systems record, and that installing software to fix a blame culture is the classic error of treating a human problem as a tooling problem. On this account a complete, timestamped, attributable record of who decided what is exactly the evidence a blame culture has always wanted and never had. Give a punitive organisation a better archive and you have not made it fairer, only more efficient at prosecution. Most of that is correct and should be conceded. A record cannot create a culture. What happens to the first competent person whose sound decision produces a bad result will do more to set behaviour in a month than any system will do in a year.

The honest position is narrower. The claim is not that a system creates the culture. It is that the absence of a record makes the culture impossible to sustain, because the commitment leaders make, we will judge the reasoning, not the result, is one they cannot keep without evidence nobody currently keeps. The record does not supply the fairness. It removes the excuse that fairness was not possible, and it makes the difference between the two kinds of organisation visible instead of arguable.

Risk	Control in Red Bear: Enterprise Decision Intelligence
The record becomes an instrument of blame	Scope: the decision is captured, not the conversation, with consent recorded and non-decision content discarded (FR-REC-04). Access: each person sees their own record and the criteria applied; oversight follows the reporting line (FR-PLT-08). Order of events: if the first use of the register is a review of something that went wrong, the organisation has answered which kind it is; review a success first
A blended score	Process and outcome scored separately; no interface combines them; dumb luck is marked the most dangerous cell (FR-OUT-05, FR-OUT-08)
Believability weighting hardens into rank	The ledger is per decision class, recommends and never applies, and carries process quality alongside outcome (FR-DEL-03)
Reconstruction contaminated by hindsight	Reasoning, confidence and dissent are immutable once written, before the outcome is known (FR-REC-12, FR-REC-08)

Risk	Control in Red Bear: Enterprise Decision Intelligence
Escalation improved, not understood	Referral patterns distinguish one case type repeatedly referred from referrals consistently authorised unchanged (FR-RTE-05)
Lessons attached to names	Pattern detection across classes with no attribution; coaching by invitation with nothing attributable kept (RBP FR-DL-04, RBP FR-DC-04)
Red Bear overstates what it has seen	Red Bear does not have evidence that any organisation has bought its way out of a blame culture, because Red Bear has no customer results; every measure here is baselined on the customer's own register (FR-PLT-05)

Sources

Every claim attributed to a source names it in the text as *author*, *work*, *section*; the full references are below. Statements without a citation are Red Bear's own position, developed for Red Bear: Enterprise Decision Intelligence. Quotations are limited to fifteen words.

Cited as	Source	Authority
Duke	A. Duke, *Thinking in Bets*, Portfolio (2018): resulting	Practitioner reference
Drucker	P. Drucker, *The Effective Decision*, Harvard Business Review (January 1967): classifying the generic and the exceptional	Academic reference
Kahneman, Sibony and Sunstein	D. Kahneman, O. Sibony and C. Sunstein, *Noise: A Flaw in Human Judgment*, Little, Brown (2021): independent judgements before discussion	Academic reference
Klein	G. Klein, *Performing a Project Premortem*, Harvard Business Review (September 2007): the pre-mortem	Academic reference
Lovallo and Sibony	D. Lovallo and O. Sibony, *The case for behavioral strategy*, McKinsey Quarterly (March 2010): the 1,048-decision study	Practitioner research
Blenko, Mankins and Rogers	M. Blenko, M. Mankins and P. Rogers, *Decide & Deliver*, Harvard Business Review Press (2010): quality, speed, yield and effort	Practitioner research
Dalio	R. Dalio, *Principles*, Simon & Schuster (2017): believability weighting	Practitioner reference
Institute of Directors	Institute of Directors, guidance on the role and responsibilities of the board: delegation of authority, not liability	Professional body
Book 1	D. Oliver, *Empower, Trust, Verify*, Book 1 (2026), chapters 2 and 7: the quality quadrant; decision classes	Red Bear
Red Bear	EDI-FSD-01 v3.0 draft; Red Bear Platform Functional Specification RBP-FSD-01	Red Bear

Glossary

Term	Meaning in this paper
Resulting	Judging a decision by its result rather than by the reasoning available at the time (Duke)
Quality quadrant	Deserved success, bad luck, dumb luck, deserved failure; decision quality against outcome quality with no combining axis
Felt safety	What people expect to happen if a decision goes wrong; measured by the escalation rate of decisions they were entitled to make
Decision class	Routine, discretionary, significant or critical; sets the evidence standard, memory standard, context depth and Cartesian threshold
Independent assessment	A judgement recorded before discussion; the spread between assessments is noise, and is information

Term	Meaning in this paper
Pre-mortem	Imagining the decision has already failed and writing down why, before the decision is made (Klein)
Believability weighting	Giving more weight to those who have repeatedly and successfully done the thing in question; per domain, visible, recomputed, never the only route for dissent (Dalio)
Earned trust ledger	The evidence about an actor's record per decision class, from which the product recommends a dial setting it cannot apply
Organisational memory	The rejected options, stated assumptions and dissent behind past decisions; not the documents that record the conclusion
Policy of the day	The rule and limits in force when the decision was evaluated, retained with the record

The series

This paper is one of ten. Each stands alone and each follows the same nine sections, so a reader who knows one knows where to look in the next.

Paper	Title	Theme
EDI-WP-01	Cartesian Decision Making	Learning what happened next
EDI-WP-02	The Cost of Organisational Indecision	The decision that never gets made
EDI-WP-03	Trust and Verify	Visibility, not control
EDI-WP-04	Delegation Without Losing Governance	The boundary and the room
EDI-WP-05	Building a Decision-Making Culture	Judge the decision, not the result (this paper)
EDI-WP-06	Decision Impact Analysis	What a decision dataset can and cannot tell you
EDI-WP-07	The Decision-Aware Enterprise	Decisions in flight
EDI-WP-08	Decision Making Anywhere	The channel is part of the governance
EDI-WP-09	When Micromanagement Is the Right Management	Oversight is a dial
EDI-WP-21	The Customer Journey	Connect, Observe, Govern, Learn, Compound

Red Bear Artificial Intelligence · Abu Dhabi, United Arab Emirates · info@redbearai.net · www.redbearai.net