

RED BEAR ARTIFICIAL INTELLIGENCE

ABU DHABI, UNITED ARAB EMIRATES

Cartesian Decision Making

Decision intelligence that learns not just what we decided, but what happened next

Every large organisation keeps meticulous records of what it did. Almost none keeps a usable record of what it decided, why, or what it expected to happen. This paper explains the discipline of holding the choice and the result apart, the four questions asked before a significant commitment, and how Red Bear: Enterprise Decision Intelligence captures the record as a by-product of how people already work.

Empower. Trust. Verify.

White paper EDI-WP-01 · Red Bear white paper series · Version 2.0

26 September 2026

Summary

The transaction survives; the reasoning evaporates. When a decision is questioned a year later, the organisation reconstructs it from memory and from whoever is still employed. When a strategy underperforms, nobody can say which assumption failed. When a manager's judgement is assessed, it is assessed on outcomes, which rewards luck and punishes sound reasoning that met an unlucky world.

Cartesian decision making is the discipline of holding two things apart that organisations routinely collapse together: the choice that was made, and the result that followed. A good outcome does not prove the decision was good. A bad outcome does not prove it was wrong. Until an organisation can tell the difference, it cannot improve its decisions; it can only react to its results.

The discipline has two halves. Looking forward, it asks four questions before a significant commitment, not one: what happens if we act, what does not happen if we act, what happens if we do not, and what does not happen if we do not. Most people answer only the first, and the other three carry the information that usually changes the decision. Laid out on two axes, act or not and happens or not, the four answers form four quadrants, which is where the name comes from. Looking backward, it compares the decision and its original assumptions against what occurred, which is only possible if the assumptions were written down at the time.

This paper explains the model, why recording the decision alone is insufficient, and what an organisation must retain for the comparison to be worth anything. It then describes how Red Bear: Enterprise Decision Intelligence (formerly GuardrAll) captures that record in the flow of work rather than as an administrative task nobody completes. The argument stands without the product. An organisation that adopted nothing but the four questions and a review date would decide better than one that did not. The product exists because, at scale, discipline that depends on individual diligence does not survive contact with a busy quarter.

WHO SHOULD READ THIS

- Chief executives, chief operating officers and chief financial officers
- Chief risk officers and heads of internal audit
- Leaders of functions that make many comparable decisions: procurement, credit, pricing, capital allocation
- Anyone asked to explain, a year on, why a decision was made

WHAT YOU WILL BE ABLE TO DO AFTERWARDS

- Distinguish decision quality from outcome quality, and say why the distinction is humane as well as analytical
- Ask the four quadrant questions of a proposal and read the answers as a shape
- Name what a decision record must hold for a retrospective to be evidence rather than storytelling
- Set the threshold above which the four questions become mandatory

1 What it is

A decision is a commitment of resources or reputation, made under uncertainty, by someone with the authority to make it. It has four parts: a question that admits more than one answer, the options that were genuinely available, the choice made and the grounds for it, and the commitment that follows. A record missing any of the four is an audit entry, not a decision record (Oliver, *Empower, Trust, Verify*, Book 1, chapter 2).

Cartesian decision making adds two instruments to that record, one used before the decision and one after.

1.1 Looking forward: the four quadrant questions

Asked to justify a course of action, people describe the upside of acting. It is the view least likely to contain anything they had not already considered. The method makes the other three views compulsory. Above a risk threshold set by the accountable director, the decider answers four questions before acting:

	It happens	It does not happen
We act	What happens if we do this?	What does not happen if we do this?
We do not act	What happens if we do not?	What does not happen if we do not?

The cost of standing still sits in the bottom-left. What acting quietly forecloses sits in the top-right. The risks avoided and options preserved by not acting sit in the bottom-right. These are the quadrants people skip when left to themselves, and between them they hold most of what a decision turns on. The name is geometric: two axes, four quadrants. It carries no attribution to Descartes.

Three rules govern the answers. The requester writes them, not the product. The pattern across the four is **read as a shape, never scored or totalled**; any interface that produces a Cartesian score contradicts the method. And the threshold is a first-class object with a standing duty to watch it for drift, defaulting to the boundary of the significant class (Book 1, chapter 9).

The device is close in spirit to the pre-mortem, in which a team imagines the decision has already failed and works backwards to explain why (Klein, *Performing a Project Premortem*, Harvard Business Review, 2007). Both make the uncomfortable views compulsory rather than optional.

1.2 Looking back: separating the decision from the outcome

Hold two variables apart. Was the decision well made? And did the outcome we were concerned with occur? Decision quality and outcome quality are scored separately, and the pair is queryable as a quadrant: a good decision with a good outcome is *deserved success*; a good decision with a poor outcome is *bad luck*; a poor decision with a good outcome is *dumb luck*, the most dangerous of the four, because it rewards a method that will fail next time; a poor decision with a poor outcome is *deserved failure*.

Judging a decision by its result is what Annie Duke calls *resulting* (Duke, *Thinking in Bets*, chapter 1). It teaches an organisation to avoid decisions rather than make them well. Separating the two requires the original reasoning to have been recorded at the time. That is the humane argument for capture as well as the analytical one: recorded context lets an organisation be fair as well as rigorous.

Key points

- Four questions on two axes, answered by the requester above a director-set threshold, read as a shape
- Decision quality and outcome quality scored apart, never blended into one number
- Both halves depend on reasoning recorded before the outcome was known

2 Why it matters

Consider what an organisation can answer about a decision taken eighteen months ago, and what it cannot. It can usually answer what happened: the supplier was appointed, the discount was granted, the product launched in March.

It can rarely answer why: what options were on the table, what evidence was in front of the people in the room, what they assumed about volumes or competitors, how confident they were, who dissented and on what grounds, what they expected the result to be and by when.

That information existed. It was in a meeting nobody minuted properly, in a thread in one person's mailbox, in a spreadsheet superseded three times since, and in the heads of people who have moved on. It was not lost through negligence. It was never captured anywhere durable, because no system asked for it and no process required it.

An organisation measures what its systems make measurable. Transactions are instrumented to the decimal place; decisions are not instrumented at all. The gap reflects what the software of the last thirty years was designed to do, not the relative importance of the two.

Consequence	How it presents
Judgement is assessed on results	Without the original reasoning, the only available measure of a manager's decision-making is whether things turned out well; people learn to choose the option easiest to defend
The same mistakes repeat	A team that gets a category of decision consistently wrong cannot discover it, because the pattern is only visible across many decisions held in a comparable form
Assumptions are never tested	Beliefs about how customers behave, how long integrations take and how competitors respond are never checked, because they were never recorded as claims
Institutional memory leaves with people	When an experienced executive departs, the organisation loses the reasoning behind every significant decision they made
Regulatory exposure rises	Supervisors increasingly ask how a decision was made, on what basis, by whom under what authority; reconstructing that is slow and unconvincing
Governance becomes opinion	Arguments about whether oversight is proportionate cannot be settled with evidence, so they are settled by seniority

An illustration. A manufacturer consolidates two suppliers into one, expecting a meaningful reduction in unit cost. Eighteen months later unit cost has risen. The executive team divides between those who believe the consolidation was a mistake and those who believe it was sound but poorly executed. Both positions are argued from the same evidence, the cost went up, because the original expectation, the assumptions behind it and the risks explicitly accepted were never written down. The discussion is settled by seniority, and the organisation learns nothing it can apply to the next consolidation.

3 The challenges

Four causes, none of which is incompetence, and a fifth that is a matter of scale.

Challenge	Why it is hard	What goes wrong without a method
Systems record transactions	An ERP records that a purchase order was approved at a value, by a user, on a date; it has no field for *why* because the reasoning is not needed to process the order	An exhaustive account of activity and a silent one about judgement
The reasoning lives in conversation	A commitment reached on a call exists only in the memory of the people on it; an approval given in an email thread sits in one mailbox	The minute records that a topic was discussed, which is true and useless

Challenge	Why it is hard	What goes wrong without a method
Nobody is asked for the expectation	Approval processes ask whether a decision is permitted, not what the decider expects to happen or by when	There is nothing to compare the outcome against, so the comparison is never made, and nobody misses the field
Hindsight rewrites the record	Reasoning reviewed after the fact is reconstructed, not retrieved, and reconstruction is contaminated by knowing the answer (Kahneman, *Thinking, Fast and Slow*, chapter 19)	A retrospective without a contemporaneous record is closer to storytelling than analysis
Discipline does not survive a quarter	The four questions are asked when the quarter is calm and skipped when it is not, which is when they matter most; the review date is diarised and missed	The record is kept diligently by three managers and not at all by thirty

4 How Red Bear: Enterprise Decision Intelligence helps

Red Bear's position is that decision capture has to be a by-product of how people already work, and that the analysis has to be proportionate: applied to decisions that matter and invisible on decisions that do not. Two design constraints follow. Analysis is threshold-triggered: the four questions become mandatory only above a risk level the organisation's own directors set, and below it a decision is captured in a couple of fields. And the model advises; people decide. A model can ask someone to answer properly, or ask a human to take a look, and nothing more.

Challenge	What Enterprise Decision Intelligence does	Requirements
Systems record transactions	Enterprise application connectors read the decision at the transaction; productivity and messaging connectors capture in Teams, Outlook and chat; voice intake extracts the decision and its spoken reason from a transcript before the transcript is discarded	FR-CAP-01, FR-CAP-02, FR-CAP-03, FR-CAP-05
Reasoning lives in conversation	Decision detection finds implicit decisions in conversation and confirms them with the owner before anything is recorded; a candidate not confirmed within the class's window is discarded, not held	FR-CAP-08, FR-CAP-02
Nobody asks for the expectation	The point-of-decision prompt asks four fixed questions at the moment of commitment, confirm-first; rationale prompting asks for the missing Cartesian elements scoped by the class's evidence standard, and a skip is recorded with who skipped	FR-CAP-12, FR-CAP-10
The four quadrants are skipped	Cartesian rationale analysis checks each record against the four quadrants, relevance and plausibility per answer, never a score, and finds which questions are most often unanswered across classes; the platform's engine holds the two axes and the quadrant plot	FR-OUT-04, RBP FR-CA-01, RBP FR-CA-02
Hindsight rewrites the record	The evidence vault is append-only and tamper-evident; audit replay reconstructs a decision as it stood, with the policy version, trust dial and context the decider saw, never re-rendered against current settings	FR-REC-08, FR-REC-09
The loop never closes	Outcome tracking links each decision to what happened afterwards, only from a connected system reporting against the decision's identifier or a review prompt answered by the accountable owner; an outcome the product guessed is not evidence	FR-OUT-01, RBP FR-DL-01

Challenge	What Enterprise Decision Intelligence does	Requirements
Judgement is assessed on results	Decision quality scoring scores the process separately from the outcome, computed without access to any measure series; the quality quadrant is queryable with no combining axis, and dumb luck is marked the most dangerous cell	FR-OUT-05, FR-OUT-08
Lessons attach to people	Patterns are 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	RBP FR-CA-05, RBP FR-DL-04, RBP FR-DC-01

Design rule

The requester writes the four answers; the product checks them. Low relevance returns the answer to the requester with a note, and this is non-blocking. Low plausibility never goes back to the requester but routes to an available authorised person with the discrepancy shown side by side. No answer is rejected by software and no judgement is overruled by a score.

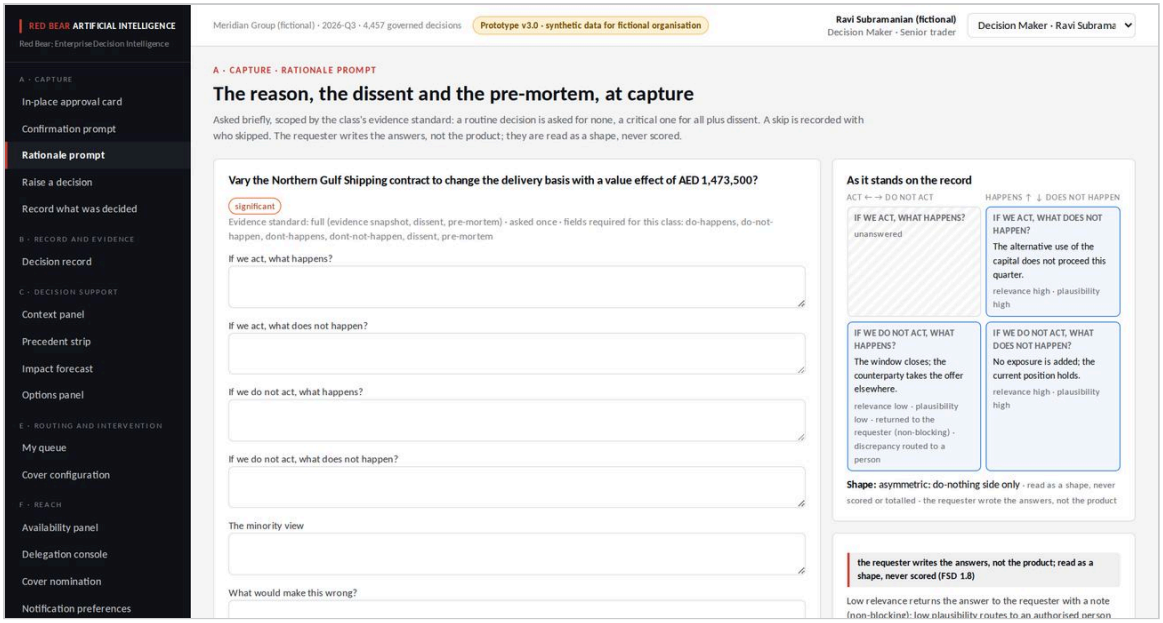
5 How it is used

Step	Who	What happens
1 Set the threshold	Accountable director, Policy Maker	The Cartesian threshold is set per decision class, defaulting to the boundary of the significant class; it is a configurable object with a duty to watch it for drift
2 Capture	Decision Maker	A decision arrives through whichever route it was made in; below the threshold it is captured in a few fields and no question is asked
3 Answer the four questions	Decision Maker	Above the threshold the requester writes the four answers and states the expected outcome and a review date before the decision completes
4 Check	Product, then a person	Relevance and plausibility are checked per answer; a low-plausibility answer is put in front of an available authorised person with the discrepancy shown
5 Decide and record	Decision Maker, Decision Approver	The choice, grounds, dissent and confidence are written to the append-only record with the policy version in force
6 Attach the outcome	Connected system or accountable owner	At the review date the outcome is reported against the decision's identifier or requested as a short review; nothing in the original record is edited
7 Learn	Class owner, Decision Approver	The quality quadrant per class shows deserved success, bad luck, dumb luck and deserved failure; a learning finding closes only when it has changed a dial, a guardrail or a standard

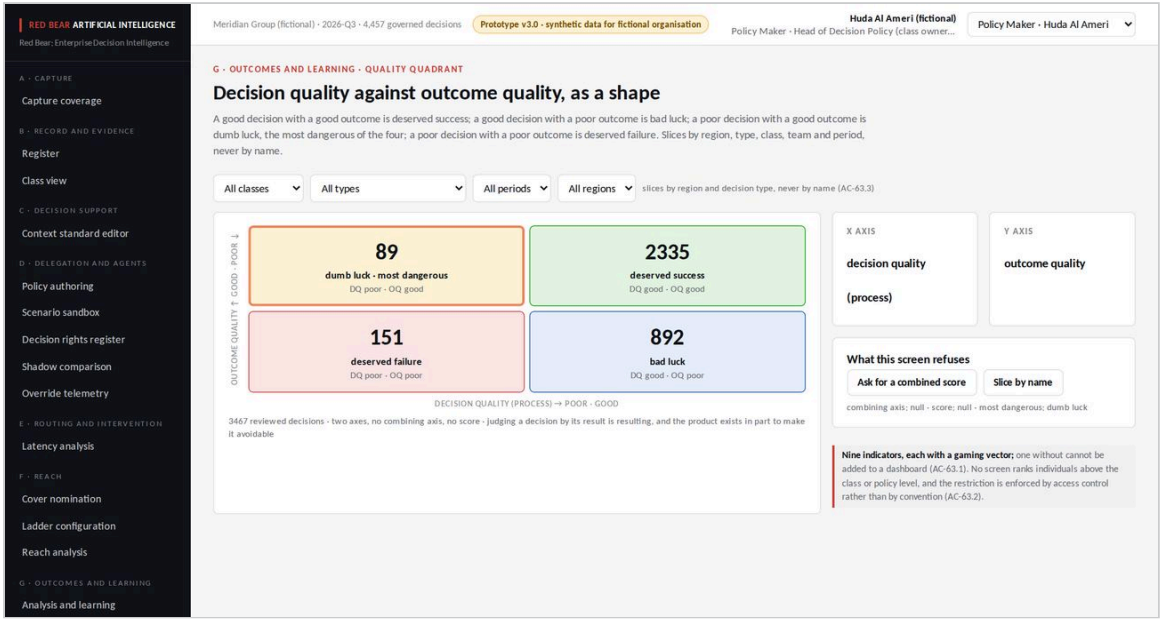
Two illustrations. A chief financial officer is asked to approve a significant systems investment. The business case is thorough on the first quadrant: efficiency gains, payback period, less manual reconciliation. The bottom-left quadrant is where the decision turns. Asked what happens if the investment is not made, the team establishes that the existing platform's support arrangement expires within the period, so declining is not a neutral option but a different and unbudgeted cost. The top-right quadrant surfaces what proceeding forecloses: the implementation consumes the same finance capability the planned acquisition integration will need. The decision changes shape before any money is committed. It is approved with a smaller initial scope and an explicit expected outcome against a date fourteen months out. Nobody has to remember to check: the record will ask.

A chief operating officer reviewing supplier consolidation decisions across the group two years later can see eleven comparable decisions, each with the reduction its sponsor expected, the assumptions stated at the time and the outcome recorded at review. A pattern emerges that no single decision revealed: consolidations expecting savings from volume leverage broadly delivered, and those expecting savings from administrative simplification broadly did not. That is not an argument about any one supplier. It is a calibration the organisation can apply to the next twenty decisions, and it was invisible until the decisions were held in a comparable form.

6 The user experience



S-03 Rationale prompt: S-03 Rationale prompt. The four quadrant questions asked above the class's threshold, with the expected outcome and review date; relevance and plausibility shown per answer, and never a score. Prototype render on synthetic data for fictional entities.



S-47 Quality quadrant: S-47 Quality quadrant. Decision quality against outcome quality for one class: deserved success, bad luck, dumb luck and deserved failure, with dumb luck marked the most dangerous cell and no combining axis. Prototype render on synthetic data for fictional entities.

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

Inputs to the Cartesian record

Input	Source	Check
The decision candidate	Productivity, messaging, enterprise-application and custom connectors, voice intake, the portal; every route terminates in one normalised candidate	Deduplicated on the decision, not the message: one commitment seen in three channels is one record with three sources (FR-CAP-07, FR-REC-02)
The four answers	Written by the requester at the point of decision	Relevance and plausibility checked per answer; low plausibility routed to a person, side by side with the discrepancy (FR-OUT-04)
Expected outcome and review date	Stated by the decider above the threshold	Completeness checked against the class's memory and evidence standards before the decision is released (FR-REC-03, FR-REC-12)
Dissent and pre-mortem	Recorded at the time for the classes whose standard requires it	Attributed and timestamped; not reconstructed later (FR-CAP-11)
Policy and dial in force	Read at evaluation time and retained with the record	Replay reconstructs against the version then in force, never today's (FR-REC-09, FR-DEL-14)
The outcome	A connected system reporting against the decision identifier, or the accountable owner's review answer	An outcome the product guessed is not evidence (FR-OUT-01)
Historic decisions	Legacy harvesting, once per engagement	Marked harvested, lower evidence standard, excluded from quality scoring (FR-CAP-09)

How it is checked

- **Connectors are read-only and consented.** Capture routes are granted by the customer administrator under least privilege, and a messaging or voice candidate is never written without confirmation from the named decision maker (RBP FR-CN-01, FR-CAP-02, FR-CAP-05).
- **Provenance travels with every claim.** Evidence is held as an immutable snapshot addressed by digest, never a live link; every observation carries its source and read time (FR-REC-12, RBP FR-PV-01).
- **The AI proposes; a named person confirms.** A detected decision, a classification below the confidence threshold, a prompted rationale: each is a proposal until a person confirms it, and an unconfirmed prompt leaves no record (FR-CAP-08, FR-REC-01, FR-CAP-12).
- **The record is never rewritten.** The comparison with the outcome is appended; the original reasoning stands as it was written (FR-REC-08).
- **Retrospective entries are marked.** A decision entered after the fact carries actual time, entry time and the difference, in every view, and is excluded from latency measures (FR-CAP-16).
- **Patterns, never individuals.** Rationale patterns are computed across classes; the quality quadrant slices by region and decision type, never by name (FR-OUT-08, RBP FR-CA-05).

8 Benefits realised

Every measure is calculated from the organisation's own register against its own baseline, established before anything is switched on. Red Bear publishes no benchmark it cannot reproduce from a live engagement (FR-PLT-05).

Benefit	How it shows up	Measured by
Reasoning survives the decision	Options, assumptions, expectation and dissent held with the choice	Decision coverage: share of in-scope decisions reaching the register with reasoning attached
Analysis before, not after	The register distinguishes answers written before the decision from answers written afterwards	Analysis timing: share of significant decisions analysed before they were taken
The loop closes	Outcomes attached at the review date without anyone remembering	Outcome closure: share of decisions reaching their review date with an outcome recorded
Fewer reversals	Decisions that considered the three skipped quadrants change shape before commitment	Reversal rate within a defined window, against baseline
Assumptions get calibrated	Which categories of assumption prove reliable and which are repeatedly optimistic	Assumption accuracy over time, per class
Fair assessment of judgement	Good decisions with bad outcomes are visible as such	Share of decisions in the bad-luck and dumb-luck cells, per class

9 Risks managed

Risk	Control in Red Bear: Enterprise Decision Intelligence
Retrospective analysis counted as deliberation	Retrospective records are marked in every view; analysis timing distinguishes the two (FR-CAP-16)
The discipline costs too much and is abandoned	Threshold-triggered: the four questions are asked only above the class's threshold, set by the customer's directors, not the vendor (FR-CAP-10)
Value expected too early	Patterns need a few hundred comparable records; coverage and timing are measurable from the first week, and no calibration figure is shown below twenty to thirty decisions (FR-REC-12)
Correlation read as causation	Rationale patterns are presented as prompts for judgement; impact is reported with confounders at equal weight and never as what the decision caused (FR-OUT-10)
The record used punitively	Quality is scored on process without access to outcomes; no interface blends the two; lessons are never attributed to a person; a person's own record is offered through private coaching, by invitation (FR-OUT-05, RBP FR-DL-04, RBP FR-DC-01)
A score replaces the shape	The four quadrants are read as a shape; any Cartesian score contradicts the method and none is produced (RBP FR-CA-02)
Hindsight contaminates review	Append-only vault and audit replay against the policy and context of the day (FR-REC-08, FR-REC-09)

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. No statistic in this paper is drawn from these sources; they are cited for the ideas named in the text.

Cited as	Source	Authority
Duke	A. Duke, *Thinking in Bets: Making Smarter Decisions When You Don't Have All the Facts*, Portfolio (2018): the concept of resulting	Practitioner reference
Klein	G. Klein, *Performing a Project Premortem*, Harvard Business Review (September 2007): the pre-mortem technique	Academic reference
Kahneman, Lovallo and Sibony	D. Kahneman, D. Lovallo and O. Sibony, *Before You Make That Big Decision...*, Harvard Business Review (June 2011): structured challenge of decision proposals	Academic reference
Kahneman	D. Kahneman, *Thinking, Fast and Slow*, Farrar, Straus and Giroux (2011): hindsight bias and the reconstruction of reasoning	Academic reference
Tetlock and Gardner	P. Tetlock and D. Gardner, *Superforecasting: The Art and Science of Prediction*, Crown (2015): recorded forecasts and calibration over time	Academic reference
Book 1	D. Oliver, *Empower, Trust, Verify*, Book 1: the Enterprise Decision Governance framework (2026), chapters 2 and 9	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
Cartesian decision making	Examining a proposed decision from four quadrants rather than one, and separating the choice made from the result that followed
Quadrant questions	What happens if we act; what does not happen if we act; what happens if we do not; what does not happen if we do not
Read as a shape	The pattern across the four answers points to the lower-risk path; it is never scored or totalled
Threshold	The risk level, set by the accountable director per class, above which the four questions are mandatory
Decision quality	Whether the process was sound: options set out, evidence to the class's standard, dissent recorded, reasoning complete
Outcome quality	What the world subsequently did, attached later and kept separate
Quality quadrant	Deserved success, bad luck, dumb luck, deserved failure; queryable, with no combining axis
Resulting	Judging a decision by whether it turned out well rather than by the reasoning available at the time (Duke)

Term	Meaning in this paper
Expected outcome and review date	What the decider believed would result, and when the comparison will be made, recorded at the time
Append-only	Records are added and superseded, never overwritten or deleted, so the earlier version remains
Pre-mortem	Imagining the decision has already failed and working backwards to explain why (Klein)
Decision coverage	The share of in-scope decisions reaching the register with their reasoning attached

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 (this paper)
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
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