RED BEAR: ENTERPRISE DECISION INTELLIGENCE · DATA SHEET · G · OUTCOMES, LEARNING AND STALLS

Find out whether your decisions were any good

A bad outcome and a bad decision are not the same thing.

KEY POINTS

- Decision quality and outcome quality are scored separately and never combined; the pair is a quadrant, and dumb luck is marked the most dangerous of the four
- An outcome is attached, never inferred; the reasoning it is judged against was written before the result was known
- Overrides are a design signal, not an audit finding; learning has exactly three destinations: a dial, a guardrail or a standard
- The non-decision is a first-class record: deferrals with a date, stalls with an owner, windows that closed

What it is

Every organisation reviews its results. Almost none reviews its decisions. When a result is bad, someone asks who approved it; when it is good, nobody asks anything at all. Both are questions about the outcome, and neither tells you whether the thinking behind it was sound.

This cluster turns the register into an asset. It attaches what actually happened to the record that predicted it; scores the process separately from how it turned out; reads overrides and outcomes as evidence about a class of decision; and gives the decisions nobody made a shape, so that they can be counted too. It is the pillar that makes the other four compound rather than merely operate, and deliberately the fourth phase of adoption, because only then is there enough history to act on.

The pain point

Uncertainty cuts both ways. A carefully reasoned decision can still produce a poor result, and a careless one can get lucky. Annie Duke's *Thinking in Bets* gives the habit of judging a decision by how it happened to turn out a name: resulting. An organisation that does it teaches its people to avoid decisions that would be visible if they went wrong, which is why approvals get escalated that nobody needed to see.

Nothing connects the decision to what happened, so the same mistakes are available indefinitely. Decided is confused with done. Reasoning quality is invisible. And the decisions nobody made are invisible too: signals arrive and nothing is decided, items age, windows close, and none of it appears on any report.

How it works

Attach the outcome, after. An outcome reaches a record only from a connected system reporting against the decision's identifier, or from a review prompt the accountable owner answers at the class's interval, one question, a short answer. An outcome the product guessed is not evidence. The commitments a decision creates are tracked, so decided and done are separate states. A decision is about something, and the product holds what: the accounts, customers, teams and suppliers it touched, with the tenant's own measures against each and a baseline fixed at the moment of decision. The impact explorer shows how each object has performed since, with confounders at the same weight; deviation is what moved, never what the decision caused, and below the confidence threshold no figure is shown.

Judge the process, not the luck. Was the evidence there to the class's standard. Were real options considered rather than one option and two decoys. Was the dissent recorded. None of those questions needs to know how it turned out, and all can be answered from the record. Lovallo and Sibony's study of 1,048 business decisions in *McKinsey Quarterly* (2010) found process mattered more than analysis by a factor of six; that is their finding, not a Red Bear result, and what it supports is a design choice: if process carries that much weight, process is the thing worth scoring. Decision quality is computed without access to any outcome measure, and no interface blends the two. They are 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 poor outcome is *deserved failure*; and a poor decision with a good outcome is *dumb luck*, marked the most dangerous of the four, because it is the one the organisation will repeat. The quadrant has no combining axis and slices by region and decision type, never by name: patterns, never people, and no league table exists.

Overrides as a design signal. A guardrail overridden once is an exception. One overridden in one decision out of five is not a control at all; it is a queue with extra steps, costing the waiting time of every decision in the class and changing nothing, because the answer at the end is almost always yes. Most governance treats a high override rate as a compliance problem with the people; this cluster treats it as a design problem with the rule. Every learning finding has one of three destinations, a dial, a guardrail or policy, or a standard, and closes only when it has changed one of them.

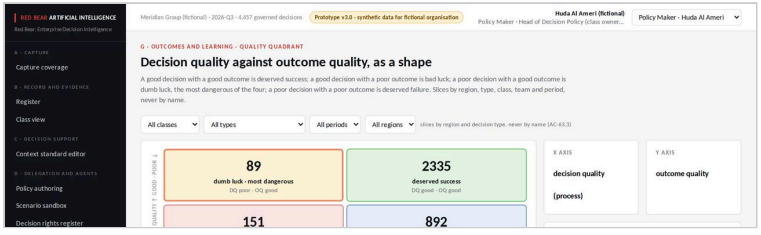
The decisions nobody made. Inaction leaves no record, so three mechanisms give an absence a shape. *A deliberate non-decision* records a choice not to act, with the reason and a date to revisit, counted alongside decisions taken. *Stall detection* finds items ageing past target with no event, signals mapped to a class with no decision, and windows that closed while a decision was open, each with an owner. *The opportunity radar* maps external signals to classes so a decision is prompted while the window is open; its output is a question to ask, not a finding to act on.

What can be inferred, and what cannot. Observing that a class of decision moves with a business outcome is not evidence that the decisions caused the outcome. A register can show a relationship; it cannot, on its own, show a cause. What this cluster contributes is the structured history such an analysis would otherwise go without. Red Bear makes no causal claim connecting governed decisions to any business result.

What you get

Capability	What it does	Requirements
Outcome tracking and execution follow-through	Outcomes from a connected system or the owner's answer, never inferred; commitments tracked so decided and done are separate	FR-OUT-01, FR-OUT-02
Decision quality scoring and the quality quadrant	Process scored apart from outcome; the four quadrants with no combining axis; nine indicators each stored with its gaming vector	FR-OUT-05, FR-OUT-08
Cartesian rationale analysis	Each record checked against the four quadrants, relevance and plausibility per answer, never a score; patterns across classes	FR-OUT-04
Impact analysis, subjects, the impact explorer and diagnostics	Effect by class on cost, revenue, risk and time, reconciled to the register; what each decision touched and how it moved, with confounders and confidence; the five diagnostics per class with trend	FR-OUT-03, FR-OUT-07, FR-OUT-09, FR-OUT-10
Stalls, non-decisions and the opportunity radar	Three stall signals with an owner; a recorded choice not to act; signals mapped to classes before the window closes	FR-OUT-11, FR-OUT-12, FR-OUT-13
Override telemetry and collective intelligence	Overrides by policy first; opt-in aggregate patterns across customers above a minimum cohort	FR-DEL-18, FR-OUT-06

Screens



S-47 Quality quadrant: The quality quadrant: decision quality against outcome quality across a type, team or period, four quadrants and no combining axis. Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

An organisation that learns: a class reviewed as a population, not through the one case that went wrong loudly; measured by outcome coverage, stalls found, and changes traced to a lesson. **Policy that changes on evidence:** thresholds move because of what happened, not what was feared. **Less defensive escalation:** separating process from result removes the reason to push a decision upwards for cover. **A dataset worth having:** the structured history any serious analysis of decision quality would require. No measure carries a target.

Platform engines

Cartesian Decision Analysis for the axes, the quadrant read as a shape and outcome known later (RBP FR-CA-01 to RBP FR-CA-05); Decision Learning for outcome reviews, lessons and patterns without attribution (RBP FR-DL-01 to RBP FR-DL-05); Decision Feedback Coaching for the private conversation on a person's own lesson (RBP FR-DC-01). Specified in the platform FSD.