RED BEAR: ALPHA ENGINE · DATA SHEET · LEARN · COLLECTIVE INTELLIGENCE

Every investigation makes the next one smarter

Predictions scored against what happened, attribution that separates the lever from the market, and, where a firm opts in, learning across the network without moving anyone's data.

KEY POINTS

- The loop closes here: the prediction made before the deal is scored against the outcome measured during ownership, and the result improves the next prediction
- Attribution is always a range: how much of the change the intervention caused, as opposed to the market, the sector or scale
- Two separate flywheels: platform-generated market intelligence for everyone, and opt-in pooled learning for firms that contribute; reciprocity, not extraction
- Nothing leaves a private environment without clearing the Safeguard Gateway, eight checks, default deny; deal-room content never crosses under any configuration
- Raw records are never centralised; only abstracted, gateway-approved patterns cross, above an aggregation threshold, and every contribution is recorded immutably

What it is

Collective Intelligence takes the prediction made before the deal and the outcome measured during ownership, scores one against the other, and uses the result to improve future predictions. Scoring alone is not enough, because a company might have grown for reasons unconnected to the intervention, so attribution estimates how much of the observed change was actually caused by the action, controlling for market movement and sector trend, and reports it as a range.

There is a second, more ambitious idea: learning across firms. If many firms contribute, the evidence base on what actually works becomes far stronger than any one firm could build. Private equity data is highly confidential, so the mechanism is deliberately restrictive. Participation is opt-in and revocable, only abstracted patterns cross, and everything passes a gateway that says no by default.

The pain point

Every deal starts from zero. The hundredth deal is modelled no better than the first, because nobody scored the first ninety-nine against what happened. The firm's own history is not a source: which levers worked, in what circumstances, at what cost and over what period is held in memory rather than in evidence. And every firm re-derives the same market picture independently, relying on benchmarks from surveys that were stale before publication, while data-sharing proposals founder on confidentiality.

Eleven pricing interventions have completed. Predicted uplift averaged 4.4%; measured averaged 3.1%; attribution suggests roughly 2.4% was caused by the intervention and the rest by market conditions. Without this pillar that comparison is never made, and the model is never adjusted for that lever in that sector.

How it works

Score, attribute, improve. When a measured outcome arrives for an initiative with a recorded prediction, the platform scores the prediction on direction, magnitude and whether the outcome fell within the stated interval, using the prediction as recorded at investment; scores are kept and never overwritten (FR-LRN-001). Attribution estimates the share of the outcome caused by the intervention rather than by market, sector or scale, as an estimate with an interval and a recorded method (FR-LRN-002). Scored outcomes and attributions retrain the uplift models in controlled, versioned releases, evaluated before release, so every estimate traces to the model version that produced it (FR-LRN-003). The result is a lever-effectiveness base by sector, size band, ownership type and cycle position, queryable from the PREDICT pillar, in which a sparse cell is reported as insufficiently evidenced rather than estimated (FR-LRN-004).

Calibration, above a minimum sample. Where decisions carry forecasts, the platform reports whether calls made at a stated probability come true at about that rate, per committee and per decider, only above a minimum sample, and only to the decider themselves and the partner group, never more widely (FR-LRN-005). The funnel view shows where names enter, where they leave and why, by stage and sector, as a pattern in the pass log rather than a league table of analysts.

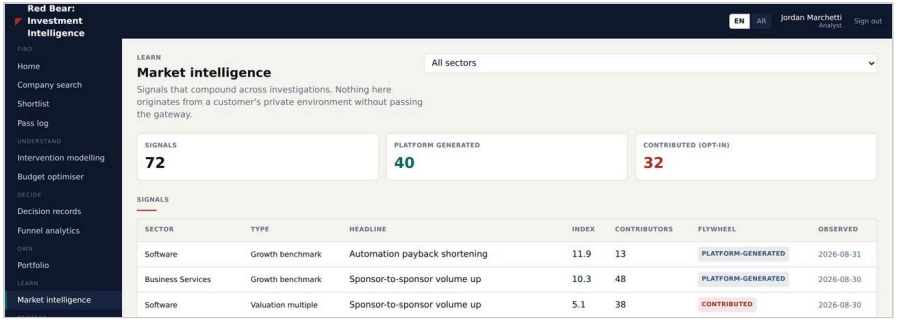
Private by default, collective by choice. A firm operates in private mode or collective mode. Private mode, the default, receives platform-generated market intelligence only; collective mode additionally contributes to and receives pooled learning. Mode is set at firm level by an administrator and changes apply prospectively (FR-LRN-006). Nothing moves by default. When a firm opts in, what leaves is never a raw record but an abstracted pattern, and it must first clear the Safeguard Gateway: proprietary customer

information, personal data, deal-room content, source licence and copyright, contractual restriction, consent and tenant policy, aggregation threshold, and re-identification risk. Every check must pass; a single failure blocks the contribution; deal-room content never passes under any configuration; and every decision is logged with its reason (FR-LRN-007). Pooled outputs are released only where enough firms have contributed that none is inferable, and sub-threshold cells are withheld rather than approximated (FR-LRN-009). Every contribution that passes is recorded immutably, with what was shared, when, under which policy version and with what aggregation, visible to the contributing firm (FR-LRN-008). A firm may revoke at any time, with no justification, and future contribution stops at once; the retention position for patterns already pooled is stated at the point of opting in (FR-LRN-010).

What you get

Capability	What it does	Requirements
Prediction scoring and attribution	Direction, magnitude and interval hit per prediction; the intervention's share of the outcome, as a range	FR-LRN-001, FR-LRN-002
Model improvement and lever-effectiveness base	Versioned retraining; effectiveness by sector, size, ownership and cycle, sparse cells declared	FR-LRN-003, FR-LRN-004
Forecast calibration	Stated probability against realised rate, per committee and decider, above a minimum sample, shared narrowly	FR-LRN-005
Participation mode and revocation	Private by default; collective by choice; revocable at any time without justification	FR-LRN-006, FR-LRN-010
Safeguard Gateway and aggregation thresholds	Eight checks, default deny, deal-room content never; pooled outputs only above threshold	FR-LRN-007, FR-LRN-009
Contribution record	An immutable record of what crossed, when, under which policy	FR-LRN-008

Screens



S-11 Market intelligence: Market intelligence: sector benchmarks from observed activity, lever-effectiveness priors, and the participation mode and gateway record for the firm.
Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

Predictions that improve with use. Each completed deal makes the next investment case sharper. **An institutional evidence base on value creation.** A defensible view of which interventions work, where, at what cost and over what period, and forecast calibration you can see. **Benchmarks that are current.** Sector comparisons reflect observed activity across a wider universe, not an eighteen-month-old survey, and patterns emerge that no single firm could see alone. **Confidentiality intact.** Deal rooms never leave the tenant, the trade is explicit, the boundary is auditable, and participation is a choice. The measures this pillar carries are **the share of completed initiatives scored** and **the interval hit rate of predictions**, both recorded by the product and not given a target here.

Platform engines

Collective Intelligence consumes the Decision Learning engine for outcome reviews, the lessons register and calibration, attributed to patterns and never to people (RBP FR-DL-01 to RBP FR-DL-04, RBP FR-DL-06), and the Privacy and aggregation service for

minimum group sizes and de-identification (RBP FR-PA-01, RBP FR-PA-03). Platform behaviour is specified in the platform functional specification, not here.