RED BEAR: ALPHA ENGINE · DATA SHEET · PREDICT · PREDICTIVE INVESTMENT INTELLIGENCE

What actually creates value, not what merely correlates with it

Model the investments under consideration, singly or combined, and see the uplift over what would have happened anyway. Or give it a budget and get a proposed package back.

KEY POINTS

- Two futures, one number: a baseline with no intervention, an invested case, and the uplift between them, which is what the investor is paid for
- Around thirty investment levers, from pricing and automation to new leadership and international expansion, each with a typical range and time to impact
- Every estimate carries its confidence interval, its method and its count of comparable cases; a thin sector gets a wide interval and the platform says so
- Combined levers overlap; the deduction is calculated and stated on screen, never presented as a simple sum
- The optimiser proposes a package to a budget; the investment director accepts, amends or discards it, and nothing is committed by the platform

What it is

If we put six million into automation and two into sales, what actually happens to this business? Predictive Investment Intelligence lets an analyst model the investments under consideration (technology, process automation, regional or international expansion, marketing, sales, a new management team), set the amounts, and see the effect on growth, revenue, income, outgoings and margin, and the enterprise value that follows. Individually, and in combination. Run the other way, it takes a total investment budget and proposes the package likely to yield the most for it.

Ordinary forecasting predicts what a company will do. An investor needs to know how much of that future is attributable to their own money and effort. A company growing 20% a year that would have grown 20% anyway represents no value created by the investor. So the pillar models two futures, and the difference between them is the number the product exists to produce.

The pain point

Value-creation plans are assertions. Without this pillar, lever selection rests on a handful of remembered successes and pattern-matching from the last portfolio company, with no control for management quality, sector or cycle, and no way to tell a good outcome from a lucky one. The uplift in the investment case is the claim an LP will test, and today it is the least evidenced figure in the paper.

With it, those calls rest on effects estimated against comparable businesses that actually made that investment, with confounders addressed explicitly and disclosed, and the uncertainty stated rather than buried. The lever is separated from the circumstances around it, and this business is analysed on its own characteristics rather than by analogy.

How it works

Baseline first. For any company the platform projects revenue, margin, income, outgoings and enterprise value over a horizon (sixty months by default) under a no-intervention assumption, from the company's own trajectory and the sector trend. The baseline is stored and versioned so every later comparison starts from the same point, and it is recomputed only on request, never silently. Baseline and invested case are always shown together (FR-PRED-001).

Forwards, then backwards. Forwards, the analyst picks levers from a taxonomy of around thirty across commercial, technology, operations, growth, people, M&A, finance and risk, each carrying a typical investment range, a typical time to impact and the outcomes it affects (FR-PRED-002). The platform estimates what each does to the outcome measures as the difference between invested case and baseline, with a confidence figure, an interval, the estimation method and the number of comparable cases behind it; uplift is never shown without its interval, and every estimate is labelled *predicted* (FR-PRED-003). An amount outside the evidenced range is scaled with diminishing returns and clamped, and the clamping is shown (FR-PRED-004). Backwards, the analyst supplies only a budget, and the platform proposes a package ranked by confidence-weighted value per unit invested, reporting what is allocated and what is left over rather than forcing money into a lever (FR-PRED-006).

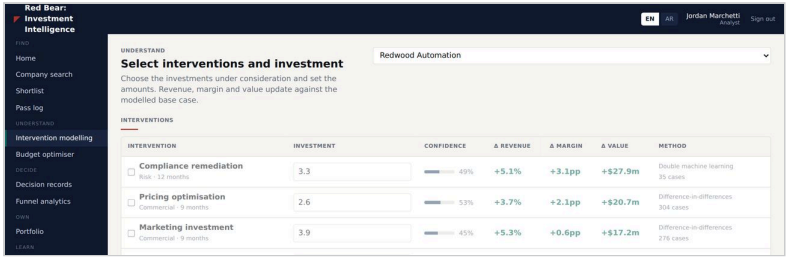
Effects, not sums. Two levers that both improve margin partly overlap. When several are selected the overlap is estimated and deducted from the combined effect, and the deduction is reported as a percentage on screen; a single lever incurs none (FR-PRED-005). Named scenarios are saved with the company and compared side by side across every measure and total investment (FR-PRED-009). Every estimate records the method applied and why it was chosen for that shape of data, in plain language, and that record travels into the IC pack (FR-PRED-008).

A prediction that is kept. When diligence returns reconciled evidence, affected estimates are marked stale and can be re-run; the pre- and post-diligence cases are both retained, and the change in confidence is reported (FR-PRED-007). A chosen package is promoted to a value creation plan carrying levers, amounts, predicted uplift and confidence, and the prediction as it stood at investment is held unaltered for comparison against what actually happens (FR-PRED-010). The platform estimates and proposes; the investment director decides what goes into the case.

What you get

Capability	What it does	Requirements
Baseline projection	Revenue, margin, income, outgoings and enterprise value with no intervention, stored and versioned	FR-PRED-001
Investment lever taxonomy	Around thirty levers with typical range, time to impact and affected outcomes	FR-PRED-002
Uplift estimation and scaling	Effect per lever with interval, method and comparable-case count; amounts scaled and clamped visibly	FR-PRED-003, FR-PRED-004
Combination and overlap	Combined effect with the overlap deducted and the deduction stated	FR-PRED-005
Budget optimisation	A proposed package to a budget, ranking basis disclosed, unallocated capital reported	FR-PRED-006
Scenarios and method transparency	Named scenarios compared side by side; method and its rationale recorded per estimate	FR-PRED-008, FR-PRED-009
Re-run and hand-off	Stale estimates re-run on new evidence with both cases kept; a package promoted to the value creation plan	FR-PRED-007, FR-PRED-010

Screens



S-06 Intervention modelling: Intervention modelling: the analyst selects levers and amounts; revenue, margin and value update against the modelled base case, with the interval and the overlap deduction shown. Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

Capital to what works. Allocation follows evidence, not anecdote: the case rests on effects drawn from comparable businesses rather than a remembered success. **Defensible attribution.** Uplift claims that survive IC scrutiny, because every figure carries its interval, its method and its evidence, and the investor's own contribution is stated separately from what the market would have done anyway. **A compounding view.** Because the prediction as it stood at investment is kept and later scored against the outcome measured during the hold (FR-OWN-004, FR-LRN-001), the firm learns which levers work

where. The measure that carries this pillar is **predicted uplift against measured outcome**, per lever and sector; it is recorded by the product from the first completed initiative and is not given a target here.

What it depends on

Models are trained first on public data (filings, registries, disclosed transactions, published financials), so the platform arrives with a view of how comparable businesses have behaved, then sharpen against your own intervention record and the outcomes the OWN pillar supplies (FR-OWN-010). Thinly evidenced sectors get wider intervals; people who know the sector challenge the estimates. Inference runs through the platform's AI gateway in the tenant's region (RBP FR-AI-02, RBP FR-AI-04); every output is labelled predicted.