

Deal Robustness Analysis

Stress-testing an investment thesis under deep uncertainty

Demonstration case: “VoltCycle BV” — synthetic battery-recycling roll-up (NL/BE)

Method: Exploratory Modeling & Analysis (EMA) with PRIM scenario discovery

Prepared as: capability demonstration for climate & impact growth investors

Important: All company data, deal terms and figures in this report are fully synthetic and for illustration only. They do not describe any real company, transaction or portfolio. The methods, however, are real, peer-reviewed, and run on open-source software (EMA Workbench, TU Delft).

Executive summary

2.98x

base-case MOIC (24.4% IRR)

2.37x

median future, 5,000 runs

15.1%

of futures underperform
($<1.5x$)

1.9%

of futures lose capital ($<1.0x$)

A conventional diligence model answers one question: *what does this deal return under our assumptions?* For VoltCycle, the answer is an attractive **2.98x MOIC / 24.4% IRR**. This analysis answers the harder question: **under which conditions does that answer stop being true?**

We ran the same deal model across **5,000 plausible futures** spanning the key uncertainties no one can reliably forecast — market growth, subsidy policy, input costs, competitive pricing pressure and exit sentiment. Three headline findings:

1. The base case is quietly optimistic. The median future returns 2.37x, roughly 20% below the deterministic figure. The thesis remains genuinely attractive — 45% of futures clear 2.5x — but the single-number view overstates the central expectation.

2. Failure is a combination, not a single risk. Scenario discovery (PRIM) shows that virtually all underperformance is concentrated where **three conditions coincide**: revenue growth below $\sim 11.5\%$, exit multiples below $\sim 7.0x$, and price compression above $\sim 2.2\%$ per year. Inside that region, 99.6% of futures fail; outside it, the deal is robust.

3. The subsidy risk everyone worries about is not the killer. A full subsidy phase-out, on its own, does not push the deal below threshold — provided growth and exit conditions hold. Diligence attention and deal protections should be weighted accordingly.

Each failure condition maps to a monitorable early-warning indicator (“tripwire”) and to concrete deal-structuring actions, detailed in Sections 5 and 6.

1. Why standard diligence models under-serve climate deals

Growth-stage climate investments sit at the intersection of technology cost curves, policy regimes, and adoption dynamics — domains where the future is **deeply uncertain**: credible experts disagree not only on the values of key variables but on their probability distributions. In that setting, a base case plus two sensitivities is not caution — it is three hand-picked points in a space of thousands.

Exploratory Modeling & Analysis (EMA) inverts the logic. Instead of asking “what is the most likely future?”, it asks “across all plausible futures, where does this decision hold up — and where exactly does it break?” The approach originates in climate policy and infrastructure planning (RAND, TU Delft) where betting on a single forecast is not an option. This report applies it to a deal.

The analytical frame (XLRM)

	Meaning	In this analysis
X	Exogenous uncertainties	Market growth, subsidy regime, input costs, price compression, exit multiple
L	Levers (decisions)	Entry price, structure, hold period (fixed here; optimised in a full engagement)
R	Relations (the model)	Revenue/margin engine linking uncertainties to exit equity value
M	Metrics	MOIC, IRR, probability of loss / underperformance

2. The synthetic deal and the model

“**VoltCycle BV**” is a fictional battery-recycling roll-up in the Benelux: €25m revenue, 14% EBITDA margin, acquired at 8.0x EV/EBITDA with 30% net debt and a 5-year hold. An operational improvement program adds +0.4pp of margin per year; 20% of revenue depends on subsidy-supported demand.

The model is deliberately a **glass box** — every equation is inspectable in the accompanying notebook. Revenue growth is dampened by subsidy phase-out (on its dependent slice) and by competitive price compression; margin gains are eroded by input-cost inflation; exit equity is exit EV minus remaining debt. The value of the method lies not in model sophistication but in how the model is interrogated.


Uncertainty ranges explored

Uncertainty	Range	What it represents
Revenue CAGR	5% – 30%	Underlying market growth for battery recycling
Subsidy support	0.0 – 1.0	Full phase-out (0) to current regime persists (1)
Input cost inflation	1% – 8% /yr	Materials and labour cost scenarios
Price compression	0% – 6% /yr	Price erosion from competitor entry
Exit multiple	5.0x – 11.0x	EV/EBITDA sentiment at the exit window

Ranges are sampled with Latin Hypercube sampling (5,000 scenarios) — no probability weighting is imposed, precisely because such weights cannot be defended under deep uncertainty. In a live engagement, ranges are set jointly with the deal team.

3. Results I — from one number to a distribution

Figure 1 — The same deal across 5,000 plausible futures




Metric	Value	Reading
Base case MOIC	2.98x (IRR 24.4%)	The deterministic “Excel answer”
Median future	2.37x (IRR 18.9%)	Central tendency across all plausible futures
P10 – P90 band	1.34x – 4.03x	80% of futures land in this range
P(MOIC ≥ 2.5x)	45.4%	The upside case is real and substantial
P(MOIC < 1.5x)	15.1%	Roughly one future in seven underperforms
P(MOIC < 1.0x)	1.9%	Outright capital loss is rare but not negligible

Interpretation. The deal is attractive but the deterministic base case sits ~0.6x above the median future — a structural optimism that arises whenever point estimates are set near the favourable end of genuinely uncertain ranges. The honest summary for an IC is not “3.0x” but “most likely 2.0–2.8x, with a 45% shot at 2.5x+ and a 15% underperformance tail that is concentrated in identifiable conditions.” Identifying those conditions is the next step.

4. Results II — what exactly breaks the deal (PRIM)

PRIM (Patient Rule Induction Method) searches the 5-dimensional uncertainty space for the smallest “box” in which failures (MOIC < 1.5x) concentrate. It reports **density** (share of futures inside the box that fail) and **coverage** (share of all failures the box captures).

Figure 2 — Where the failures live: scenario discovery (PRIM)




The discovered failure region

Condition (all three together)	Threshold	Share of range
Revenue CAGR	below ~11.5%	bottom quarter of the growth range
Exit multiple	below ~7.0x	bottom third of the exit range
Price compression	above ~2.2% /yr	upper two-thirds of competitive pressure

Box quality: 99.6% density, 36% coverage. Inside this region, virtually every future underperforms; the single box explains over a third of all failures (additional boxes can be peeled to map the remainder). Failure requires the *coincidence* of slow growth, a weak exit window and sustained pricing pressure — no single risk does it alone.

The most valuable finding is what is absent. Subsidy support does not appear in the failure box:

Figure 3 — Subsidy phase-out alone does not break the deal



Failures (navy) occur across the entire subsidy axis — a full phase-out, by itself, does not break the deal, because only 20% of revenue is subsidy-dependent and the ops program offsets the drag. This reallocates diligence attention: the intuitive headline risk (policy) matters less than the mundane trio of growth delivery, exit timing and pricing power. A three-column sensitivity table cannot produce this kind of finding.

5. Actionable recommendations

5.1 At the investment decision (pre-signing)

#	Action	Rationale from the simulation
1	Underwrite at the median (2.4x), not the base case (3.0x); treat the gap as required margin of safety in entry price negotiation.	Base case sits ~20% above the median of 5,000 futures.
2	Negotiate downside protection tied to growth delivery — e.g. an earn-out or ratchet triggered if revenue growth runs below ~11–12%.	Growth below ~11.5% is the strongest single failure condition (present in the failure box).
3	Preserve exit-timing flexibility: avoid structures forcing exit in a fixed year; secure board rights to extend the hold if sector comps sit below ~7x.	Exit multiples below ~7.0x define the second failure condition; timing flexibility is the cheapest hedge.
4	Make pricing power a diligence deep-dive: contract structures, switching costs, competitor pipeline.	Sustained price compression above ~2.2%/yr is the third failure condition — and the one management can influence most.
5	De-prioritise subsidy scenario debates in IC; document the exposure (20% of revenue) and move on.	Subsidy phase-out alone does not breach the return threshold in any discovered failure box.

5.2 After the investment — the tripwire monitoring sheet

Each failure condition converts into an early-warning indicator that can be tracked with data the fund already receives:

Tripwire	Trigger (from PRIM)	Data source	Response if triggered
Growth tracking	Trailing 12m revenue growth trending < ~11.5%	Quarterly management accounts vs. plan	Accelerate commercial plan; revisit buy-and-build pace
Exit-window sentiment	Sector EV/EBITDA comps drifting < ~7.0x	Listed comps, recent transaction multiples	Prepare hold-extension case; delay exit prep spend
Pricing power	Realised price erosion > ~2.2% /yr	Price/volume bridge in monthly reporting	Contract renegotiation program; differentiation investment

Decision rule: if two of the three tripwires fire simultaneously, the deal has entered the discovered failure region. That is the moment to act — while options (ops acceleration, exit re-timing, pricing intervention) still exist — rather than discovering the outcome at exit. This converts the analysis from a one-off diligence exhibit into a living risk-management instrument reviewed each quarter.

6. What can we do for you?

Everything in this report was produced from a deliberately simple synthetic model. In an engagement, the identical pipeline runs on **your deal, your operating model, your assumptions** — and the levers you actually control (entry price, structure, reserves, hold period) are optimised for robustness rather than held fixed.

What we need from you

Input	Format	Effort on your side
The deal's operating model	Excel as-is (we wrap or rebuild it — no rework needed)	Send the file
Key uncertainties & plausible ranges	1–2 hour working session with the deal team	One meeting
Definition of “failure” and “success”	Return thresholds (e.g. MOIC < 1.5x, IRR < hurdle)	One decision
Candidate deal structures (optional)	Term-sheet variants to compare (ratchets, earn-outs, reserves)	Existing documents

What the model returns

Deliverable	What it gives your IC
Outcome distribution	Full MOIC/IRR distribution with P(loss), P(underperformance), P10–P90 band — replacing the three-column sensitivity page
Failure map (PRIM)	The specific, quantified combinations of conditions under which the deal breaks — and which feared risks turn out not to matter
Sensitivity ranking (Sobol)	Which uncertainties drive outcome variance, ranked — focusing diligence effort where it pays
Structure comparison	Candidate term-sheet variants scored on robustness across all futures, not just the base case
Tripwire sheet	Quarterly monitoring indicators with thresholds and pre-agreed responses, ready for portfolio reviews
Reproducible notebook	The full analysis as auditable, re-runnable code — update it each quarter as actuals come in

Engagement formats

Format	Scope	Timeline
Pilot: single-deal stress test	One live or recent deal, your model, one workshop → this report on your numbers	2–3 weeks
Portfolio company robustness review	Growth plan of one portfolio company tested across futures; strategy recommendations	3–4 weeks
Portfolio vulnerability scan	Correlated macro/policy conditions that hit multiple holdings at once — LP-report ready	4–6 weeks

Methods: Exploratory Modeling & Analysis (Bankes 1993; Kwakkel & Pruyt 2013); scenario discovery with PRIM (Bryant & Lempert 2010); implemented in the open-source EMA Workbench (TU Delft). All software is open source — you are never locked into a proprietary tool.

Disclaimer: This document is a methodology demonstration using entirely synthetic data. It does not constitute investment advice or a recommendation regarding any security or transaction.