INSPIRE REPORT

5 Steps to Make *Your Innovations Bankable*

Defining the justifiable development spend required to take a technology from not yet feasible to investable: a five-model method.

Arij van Berkel, Ph.D.
Chief Product Officer

Lux Research

Why good innovations die early

Most innovation decisions are made with a tool built for the wrong question — and the best ideas are the ones it rejects first.

The conventional gate for an innovation investment is a deterministic net-present-value (NPV) calculation: estimate the costs, estimate the revenues, discount them, and approve the project if the number is positive. It is a sound test for a known asset. It is a poor test for an innovation.

The problem is structural. A first-of-a-kind technology is, by definition, not yet at the cost and performance it will reach after development. Run today's numbers and the verdict is almost always the same: a cost base above the market price, a return well below the hurdle rate, and no investment. The gate does its job and the option closes — not because the innovation is bad but because it was measured before it was developed.

This is the central dilemma every R&D and innovation leader faces. You are asked to commit resources to an idea whose business case cannot yet be demonstrated against a financial standard designed for ideas whose business cases already exist. Used naively, the standard systematically destroys the options most worth holding.

The escape is not to abandon financial rigor; it is to apply the right rigor. An innovation program does not commit a company to build anything. It buys the right to build later, if and when the technology and the market justify it. That is an option, and options have value precisely because of the uncertainty that a conventional deterministic NPV analysis treats as a reason to reject the investment.

This report presents the Lux Innovation Value Framework: a sequence of five linked models for transforming an unproven idea into a funded, risk-justified, target-driven development program. Each stage answers a question that leadership teams routinely ask by applying one model, with each stage building on the last. To illustrate the method, we follow a single worked example throughout: a first-of-a-kind e-methanol production program.

Innovation is not a smaller version of a deployment decision. It is a different decision, and it deserves a different model.

A note on scope: This is a project-level method. It values and shapes one innovation project at a time; it is neither a portfolio model nor an innovation strategy. Its outputs — fund this project, at this budget, toward these targets — are inputs to portfolio and strategic decisions, not substitutes for them. The framework provides a quantified basis for those decisions by showing what a project is worth and what it will take to make it investable.

Executive Summary

From “is the case profitable?” to “what should we develop, and how much should we spend?”

Innovation decisions are often reduced to a single question asked too early: Is this profitable today? In reality, investment decisions are sequential. Should we fund the option at all? How much is that option worth? What return must the eventual deployment earn to justify its risk? And what, precisely, must development achieve to make the investment case credible?

The Lux Innovation Value Framework answers these questions in sequence using five stages, each applying a different model. The Technoeconomic Model establishes the deterministic baseline. The Risk–Reward Model extends it with Monte Carlo simulation to reveal the full distribution of outcomes, exposing opportunities that a single base case obscures. The Real Options Model values the right to invest later, defining a defensible upper bound on development spending. The Investment Hurdles Model combines a systematic risk resilience assessment with a Sharpe-ratio hurdle to test whether the eventual deployment meets the company’s risk-adjusted return requirements. Finally, the Innovation Model computes the lowest-cost set of development targets required to make a marginal project investable — and the development budget justified to achieve them.

Together, the five stages transform innovation from a binary investment decision into a quantified, staged, and actively managed process.

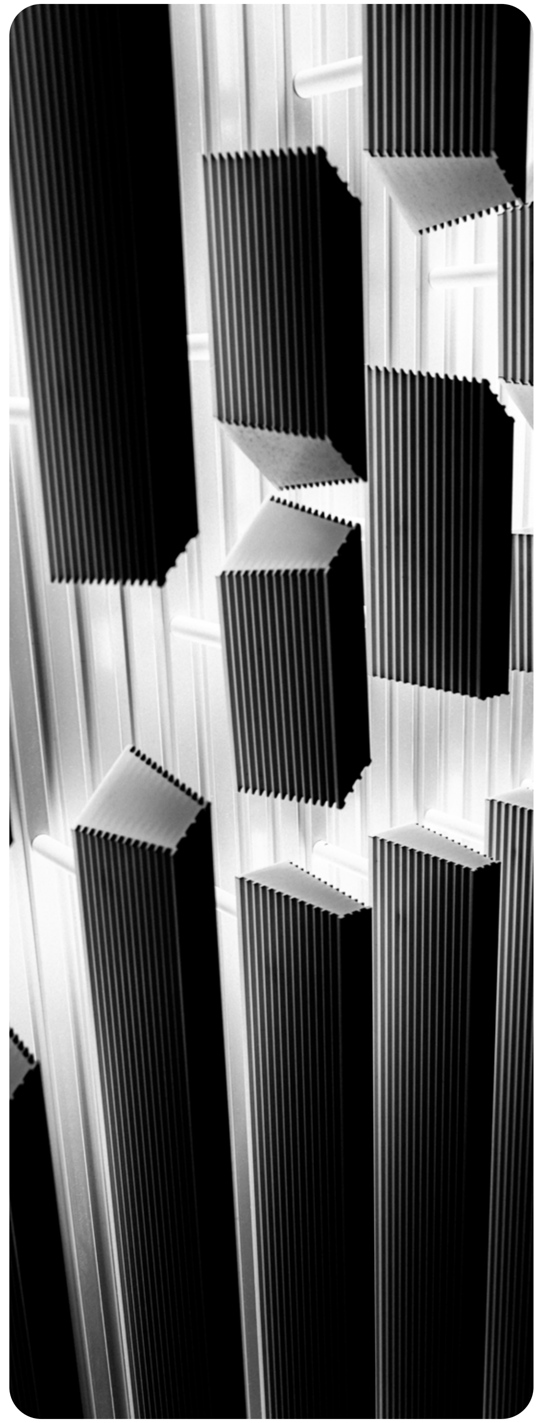

This report addresses the question leadership actually asks:

1. Even if the business case fails today, is there a realistic chance this innovation succeeds, and is that opportunity worth pursuing?

2. How much can we justifiably invest in developing this innovation, and what return must its eventual deployment earn to compensate for its risk?

3. What is the lowest-cost set of development targets that turns a marginal project into an investable one?

Five stages, one recipe

The framework consists of five stages. Each stage applies a different model to answer a specific executive question. Together, the stages form a practical workflow: Each requires a single definition and produces the answer needed for the next.

The framework follows the natural sequence of an innovation decision. It begins by asking whether the business case works as it stands. It then asks how uncertainty changes that verdict, and, crucially, whether uncertainty creates opportunity as well as risk. With the resulting distribution in hand, two questions can be answered in parallel: How much the option to develop is worth (Real Options), and whether the eventual deployment clears the company's risk-adjusted return hurdle (Investment Hurdles). The final model converts those results into a development plan by identifying the lowest-cost targets and budget required to make the project investable.

Viewed as a workflow, the five steps are:

1) define the investment; 2) define its uncertainties; 3) define the value of the development option; 4) define the required risk-adjusted return; and 5) define the development targets and budget. The table below pairs each step with the model that performs it and the question it answers.

THE INNOVATION VALUE FRAMEWORK

STAGE 1	STAGE 2	STAGE 3	STAGE 4	STAGE 5
Technoeconomic	*Risk–Reward*	*Real Options*	*Investment Hurdles*	*Innovation*
DEFINE THE STEP	**DEFINE THE STEP**	**DEFINE THE STEP**	**DEFINE THE STEP**	**DEFINE THE STEP**
The capex, opex, and production of the final investment	The uncertainties in the project's parameters	The investment option to be valued	The risk you'll accept and the reward you require	The development targets and budget
THE QUESTION IT ANSWERS	**THE QUESTION IT ANSWERS**	**THE QUESTION IT ANSWERS**	**THE QUESTION IT ANSWERS**	**THE QUESTION IT ANSWERS**
Is the base case profitable as it stands?	What is the full range of outcomes — and the upside?	How much can we spend developing it?	Does the return justify the risk?	What makes it investable and at what budget?

A first-of-a-kind e-methanol program

We carry a single worked example through all five stages: a program to develop a fleet of 10 first-of-a-kind e-methanol plants.

The inputs below are the only information provided to the models; every result in the stages that follow is computed from them. Read on its own, this is not a promising project: The capital is heavy, the product price is thin, and the cost and output are uncertain. That is exactly the profile a deterministic gate rejects — and exactly where the framework does its work. Each stage that follows applies one model to these inputs and reports what it returns.

Plant economics (per plant)	
Capex	\$505 million
Opex	\$47 million/y
Product price (e-methanol)	\$630/tonne
Production	100,000 tonne/y
Build time	3 y
Project life	28 y (25 y operating)
Program	
Plants deployed (units)	10
Scale factor/scaling exponent	1.0/0.66
Financial parameters	
Weighted average cost of capital (WACC; discount rate)	7%
General inflation	2%
Product-price inflation	2.5%
Learning & degradation	
Asset (capex) learning rate	15% per doubling
Operating learning rate	5% per doubling (on 30% of opex)
Output degradation	0.5%/y
Input uncertainty (2σ band)	
Capex	±30%
Opex	±35%
Product price	±12%
Production	±20%
Option & risk parameters	
Time to final investment decision (FID)	4 y
Risk-free rate	2.5%
Illustrative development cost	\$50 million
Company beta	1.1
Market volatility	12%

Technoeconomic

Define the capex, opex, and production of your final investment.

THE QUESTION IT ANSWERS

Is the business case profitable as it stands today?

The first model is the foundation the other four are built on. It takes a single, fully specified asset — its capex, opex, output, lifetime, learning curves, degradation, and the price its product commands — and computes the deterministic economics: the levelized cost of production (LCOP; the breakeven price), the internal rate of return (IRR; the discount rate at which that breakeven equals the actual price), and a simplified NPV.

The decisive output is the relationship between the lowest achievable cost of production and the product price. A meaningful IRR appears only when the minimum levelized cost sits below the price. For a mature asset, this is routine. For a first-of-a-kind innovation, the return often lands well below the capex — and that shortfall is the honest starting point, not the end of the analysis.

THE MATHEMATICS

$$\mathrm{LCOP}(r) = \frac{C_{\mathrm{CAPEX}} + C_{\mathrm{OPEX}}}{Q}$$

$$c_d = \frac{1+\pi}{1+r} \qquad p_d = \frac{(1-\delta)(1+\pi_p)}{1+r}$$

$$\mathrm{IRR} \quad r^* : \quad \mathrm{LCOP}(r^*) = P_{\mathrm{price}}$$

$$L(k_1, k_2, R) = \frac{1}{k_2 - k_1} \sum_{j=k_1+1}^{k_2} (1-R)^{\log_2 j}$$

The LCOP is the breakeven price; the IRR is the discount rate r at which it meets the actual price P. Discount factors carry general inflation (π), product-price inflation (πp) and degradation (δ); the Wright learning curve lowers capex and opex by a fixed fraction R for every doubling of cumulative output.*

WHAT IT GIVES YOU

A defensible, transparent base case: the cost structure, the breakeven price, the return, and the learning-curve and scale effects that will move them. It is the reference every later stage perturbs and improves upon.

WHAT TO WATCH FOR

A weak verdict here is information, not a stop sign. Treating the deterministic NPV as the decision is exactly the error the framework exists to correct. The base case is a single point on a distribution you have not yet drawn.

THE WORKED EXAMPLE

At USD 505 million of capex per plant and a USD 630/tonne product price, the fleet's LCOP lands at USD 679/tonne at the 7% WACC — about 8% above the price. The base-case IRR is just +4.9%, well short of the 7% hurdle. On a deterministic gate, the program destroys value and would be rejected outright. We carry it forward.

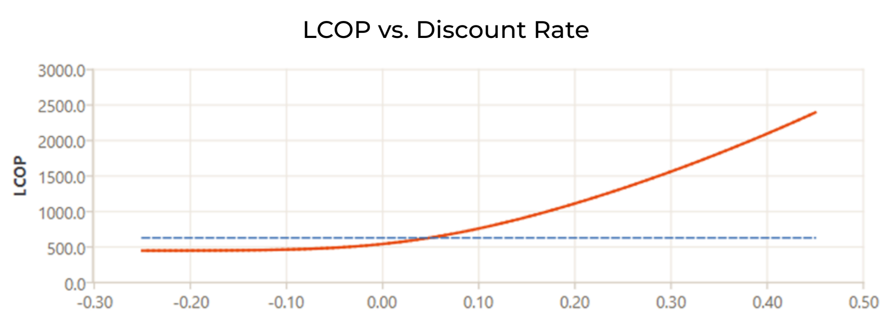

Figure 1. LCOP as a function of the rate of return as calculated by Lux Research's model software.

Risk–Reward

Define the uncertainties in the project's parameters.

THE QUESTION IT ANSWERS

What is the full range of outcomes — and is there upside worth chasing?

A deterministic base case answers the wrong question for an innovation. The inputs — capex, opex, achievable price, production, and learning rates — are not known; they are estimates with uncertainty around them. The Risk–Reward Model takes that uncertainty seriously. It wraps the technoeconomic engine in a Monte Carlo simulation, drawing thousands of scenarios in which each uncertain input varies around its expected value according to a declared risk band and accumulates the resulting distribution of returns.

This is the stage that reframes the decision. A base case can be marginal while a meaningful share of scenarios is strongly profitable — and that is precisely the situation that justifies development. The model handles this:

THE MATHEMATICS

$$\sigma_\theta = \frac{\rho_\theta\,\mu_\theta}{2} \qquad \theta_{\text{sample}} = |z+1|\,\mu_\theta$$

$$\text{accept iff} \quad \min_r \ \mathrm{LCOP}(r;\theta) \le P_{\text{price}}$$

$$\frac{1.96\,\hat\sigma/\sqrt{n}}{\max(|\bar r|,\ 10^{-4})} < \varepsilon_{\text{rel}}$$

Each uncertain input θ is sampled around its mean μ with a width set by the risk band ρ (a 40% band gives a 2σ spread of ±40%). Draws with no viable return are rejected; the mean and spread of IRR accumulate with Welford's numerically stable recurrence until the 95% confidence interval is tight enough. The model uses a Gaussian probability distribution by default; the user can change that to a selection of alternative distributions if desired.

Scenarios in which no viable return exists are rejected, and a feasibility precheck determines up front whether any favorable combination of inputs could succeed. If even the most favorable envelope fails, the model says so plainly. If it passes, the distribution shows how much of the probability mass lands in profitable territory.

WHAT IT GIVES YOU

The mean return, the 5th and 95th percentiles, and the probability of a viable case. It quantifies potential — the reward that exists if development moves the inputs in your favor.

WHAT TO WATCH FOR

The risk bands are an input, not a detail. They encode an honest view of how uncertain each estimate is, and they drive everything downstream: Option value, beta, and the development targets all flow from them. Note the incentive this creates: Because wider bands raise the option value, the bands must reflect genuine uncertainty, not a wish for a bigger development budget. They belong under the same governance as any other key assumption.

THE WORKED EXAMPLE

Across 215,000 Monte Carlo draws, 95% produce a viable return. The mean IRR is +4.6% — still short of the 7% WACC — but the distribution runs from −5.2% at the 5th percentile to +12.4% at the 95th, and 33% of viable scenarios clear the WACC. A real, sizeable case is hiding inside the uncertainty. The question becomes whether the option to capture it is worth funding.

Cumulative Probability Distribution of Return

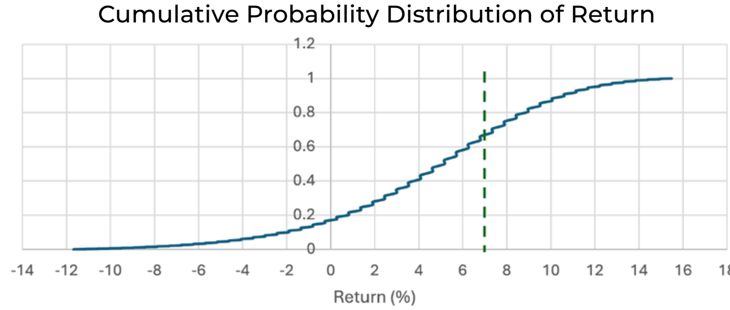

Figure 2. The cumulative distribution of the IRR as calculated by the model.

Real Options

Define the development investments.

THE QUESTION IT ANSWERS

How much can we reasonably spend developing this innovation?

A development program does not commit the company to building anything. It buys the right to invest the capital later, at an FID, once the technology and the market are clearer. If the project turns out infeasible at that point, the firm simply walks away, losing only the development cost, not the capital. That asymmetry is the entire reason innovation under uncertainty can be worth funding, even when the expected case is marginal.

The Real Options Model values the right to invest later. It takes the distribution from Stage 2 and applies the Black-Scholes European call formula, treating the present value of the project's cash flows as the underlying asset, the deployment capital as the strike, and the volatility of the project's net value, drawn directly from the Monte Carlo, as the source of option value. The output is the maximum amount worth spending on the development program. Spend less and you create value; spend more and you destroy it.

Leadership also controls the scale of the eventual deployment. A larger program increases both the option value and the capital at risk, so the tradeoff is not whether to invest, but how large an opportunity to pursue. The Real Options Model evaluates each candidate deployment scale separately, allowing leadership to choose the level of exposure it is willing to accept.

THE MATHEMATICS

$$C = S\,\Phi(d_1) - X\,e^{-rT}\,\Phi(d_2)$$

$$d_1 = \frac{\ln(S/X) + rT}{\sigma\sqrt{T}} + \tfrac{1}{2}\sigma\sqrt{T} \qquad d_2 = d_1 - \sigma\sqrt{T}$$

$$S = \text{Units} \cdot \frac{\overline{\text{Cashflow PV}}}{(1+i)^T} \qquad X = \text{Units} \cdot C_{\text{CAPEX}}$$

$$\sigma\sqrt{T} = \frac{\text{std(Net PV)}}{\overline{\text{Cashflow PV}}}$$

The Real Options Model applies the Black-Scholes European call formula. Here, S is the project's value today, X the capital required for deployment, σ√T the volatility derived from the Monte Carlo simulation, T the time to the FID, and Φ the cumulative normal distribution. The call value, C, is the maximum justifiable development spend; after subtracting the planned development cost, it becomes the net option value. Because both S and X scale with the size of the eventual deployment, the model evaluates each candidate deployment scale separately, showing how option value — and capital exposure — change together.

WHAT IT GIVES YOU

A breakeven development budget: the ceiling on what the option to deploy is worth today. Subtract your actual planned development cost and you have the net option value: Positive means the program is worth funding.

WHAT TO WATCH FOR

Time matters asymmetrically. Faster development raises the asset value more than it costs in deferred capital, so accelerating a program carries an implicit option-value benefit beyond its direct spend.

THE WORKED EXAMPLE

The fleet's present asset value is USD 2.97 billion against USD 3.59 billion of deployment capital, with a volatility ($\sigma\sqrt{T}$) of 0.55 over the four-year run-up to FID. The Black-Scholes option value is USD 195 million: the most worth spending on development. Against an illustrative USD 50 million development budget, the net option value is over USD 145 million; the option is clearly worth holding. Scale sets the magnitude: This 10-plant commitment carries USD 195 million of option value against USD 3.6 billion of capital, whereas a 20-plant commitment would roughly double the option value to around USD 390 million and double the capital at stake with it.

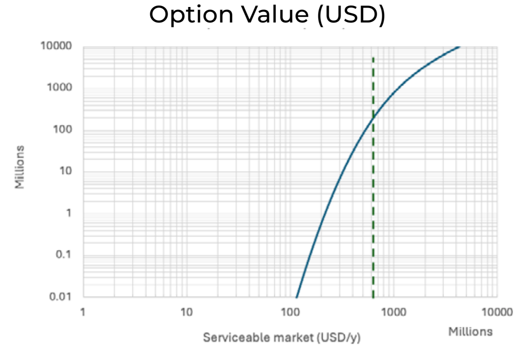

Figure 3. The option value depends on the size of the market the project can service (once deployed). The dotted line shows the scale of the project being considered, which yields an option value of USD 195 million. A bigger project would result in a higher option value but would also require more capital for the FID.

Investment Hurdles

Define how much risk you're willing to take, and at what reward.

THE QUESTION IT ANSWERS

What return must the deployment earn to compensate for its risk?

The Real Options Model asks whether the development program is worth funding. The Investment Hurdles Model asks the complementary question: If the option is exercised, will the eventual deployment earn a return commensurate with its risk? It applies two tests. Both draw on the Capital Asset Pricing Model (CAPM), but the decisive hurdle is not the standard CAPM comparison. Instead, it is a Sharpe-ratio hurdle that sets the required return according to the project's total volatility. A high expected return is not enough; it must be high enough for the risk the project carries.

The Investment Hurdles Model derives the project's risk profile directly from the Monte Carlo simulation and applies two tests. The first is a CAPM resilience test, which compares the project's systematic risk — the component that moves with the broader market — with the company's own beta. A project with a lower beta reduces the firm's market exposure. The second, and usually the binding test, is a Sharpe-ratio reward hurdle. It requires the expected return to exceed a threshold set by the project's total volatility, not just its systematic risk. Charging for total risk is a deliberate departure from textbook CAPM, which prices only systematic risk, but it is the appropriate test for an innovation held as a single, concentrated investment that the firm cannot diversify away. Together, the two tests define a project's position on a four-quadrant risk-return map.

THE MATHEMATICS

$$\beta = \sum_i m_i \frac{\mathrm{Cov}(\mathrm{IRR}, x_i)}{\sigma_i\, \sigma_{m,\mathrm{eff}}} \qquad \sigma_{m,\mathrm{eff}} = \frac{\sigma_m}{\sqrt{T}}$$

$$\mathrm{ERP} = \frac{\mathrm{WACC} - r_f}{\beta_{\mathrm{company}}}$$

$$\alpha_{\mathrm{CAPM}} = \overline{\mathrm{IRR}} - [\, r_f + \beta \cdot \mathrm{ERP}\,]$$

$$\alpha_{\mathrm{CML}} = \overline{\mathrm{IRR}} - \left[r_f + \frac{\mathrm{ERP}}{\sigma_{m,\mathrm{eff}}}\, \sigma_{\mathrm{IRR}} \right]$$

Beta is built from how each risk driver x_i co-moves with the market (weight m_i), de-annualized to the horizon T. The resilience test (α_{CAPM}) is a true CAPM comparison against the company's beta. The reward test (α_{CML}) is a Sharpe-ratio hurdle: the term ERP ⁄ $\sigma_{m,eff}$ is the market's return per unit of risk, applied to the project's own total volatility σ_{irr} — so it charges for total, not just systematic, risk, and it is usually the test that binds. A project passes only when both alphas are positive.

WHAT IT GIVES YOU

A risk-adjusted hurdle expressed in the company's own terms and a clear read on whether the deployment improves or worsens the firm's risk position. It turns "the IRR looks fine" into "the IRR is sufficient for the risk."

WHAT TO WATCH FOR

A project whose revenues and costs both track the market is partially self-hedged and carries low systematic risk — the same reason an integrated producer is less exposed than a pure-play one. Resilience can come from structure, not just margin.

THE WORKED EXAMPLE

Because its revenues and costs partly move together, its project beta is −0.24 against a company beta of 1.1; it would lower the firm's market exposure, so the resilience test passes with room to spare. But the reward test fails: The +4.6% mean return falls short of the 9% that the program's own volatility demands. The project has a strong risk profile, but its expected return is not yet sufficient.

Innovation

Define the targets and budget for the development project.

THE QUESTION IT ANSWERS

What targets make the case investable — and at what budget?

The first four stages diagnose. The fifth prescribes. For a project that is not yet investable but has positive option value, the Innovation Model asks the planning question directly: What is the smallest, least-cost set of technoeconomic improvements needed to make the project investable? It treats each improvable parameter as a development lever — lower capex, lower opex, higher output, faster learning, or reduced uncertainty — and identifies the combination that reaches investability with the lowest total development difficulty.

Crucially, the model does not optimize the project to perfection. It aims only to reach the investability boundary; every improvement beyond that point is development spend without justification. Each lever is weighted by its relative R&D difficulty, an expert input provided by the technical team, so the model identifies the easiest path rather than simply the smallest change. The result is a development assignment: a target for each lever and a maximum justifiable development budget. That budget is a ceiling, not a goal, since spending the full amount would leave no surplus.

The model stops short of the FID. It returns the conditions under which the project becomes investable — the targets and the justified budget — not a go/no-go recommendation. Whether those targets are technically achievable remains an expert judgment. If the lowest cost path requires halving opex for a budget of half a million dollars, the arithmetic may be encouraging, but the engineering team will almost certainly reject it as unrealistic. The framework's role is to make that tradeoff explicit and quantified; the decision remains with those responsible for delivering the technology.

THE MATHEMATICS

$$\min_{\delta} \sum_i w_i \, \delta_i^2$$

$$\text{subject to:} \quad \alpha(p_0(1+\delta)) \geq 0 \quad V_{\text{opt}}(p_0(1+\delta)) \geq x$$

The optimization finds the weighted set of lever changes, δ, that reaches investability at the lowest total development difficulty. Here, w_i is the relative R&D difficulty of lever i, so the objective steers development toward the easiest path. The constraints require the developed end-state to satisfy the risk-adjusted return criterion (α ≥ 0) and the affordability criterion (option value ≥ x). The justified development budget is the option value at that end-state.

WHAT IT GIVES YOU

A prioritized, quantified set of development targets and a justified maximum budget — translating financial analysis into an engineering and research agenda the innovation organization can own.

WHAT TO WATCH FOR

A result that the levers cannot close the gap is itself a finding. It says that, given their difficulty and the irreducible uncertainty, the available development paths do not reach investability — a far more useful answer than a vague "needs more work." And when it does return a path, it is stating conditions, not granting permission: Whether the targets are achievable within the budget is a call only the experts can make.

THE WORKED EXAMPLE

Holding capital cost and risk fixed, the cheapest route to investability is a 9% lift in output, while at the same time reducing opex by 7%. That raises the mean IRR to +8.1% and turns both tests positive. The most that end-state could justify spending on development is USD 430 million: That is a ceiling, not a target. Whether increasing yield by 9% and simultaneously reducing opex by 7% is actually reachable, and at what cost, is the call the engineering team now has to make.

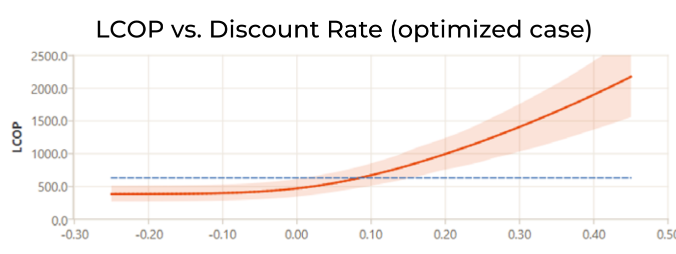

Figure 4. The levelized cost curve of the optimized case, that is, when the development project reaches its goals.

One project, five verdicts

The same e-methanol program a deterministic gate rejects emerges, five stages later, as a fundable development with a target and a budget.

Read in sequence, the five stages convert an unpromising base case into an actionable decision. Each row below is computed from the single set of inputs on the worked-example page — nothing is added along the way except the question each model is built to answer.

THE INNOVATION VALUE FRAMEWORK

STAGE 1	STAGE 2	STAGE 3	STAGE 4	STAGE 5
Define the investment	**Define the uncertainties**	**Define the project**	**Define risk & reward**	**Define targets & budget**
Technoeconomic	*Risk–Reward*	*Real Options*	*Investment Hurdles*	*Innovation*
WHAT THE MODEL RETURNS	**WHAT THE MODEL RETURNS**	**WHAT THE MODEL RETURNS**	**WHAT THE MODEL RETURNS**	**WHAT THE MODEL RETURNS**
LCOP USD 679/tonne vs. USD630/tonne price (+8%); base-case IRR +4.9%, below the 7% WACC	Mean IRR +4.6%; P5 −5.2%, P95 +12.4%; 33% of viable scenarios beat the WACC	Option value USD 195 million against a USD 50 million development cost → net +USD 145 million	Project beta −0.24 vs. company 1.1 (resilience passes); reward test short by 4.4%	+9% output and −7% opex → mean IRR +8.1%; both tests pass; justified budget USD 430 million
READ	**READ**	**READ**	**READ**	**READ**
Value destroying as it stands	**Real upside exists**	**Worth developing**	**Strong risk profile; return not yet sufficient**	**A clear, funded target**

The deterministic gate would have stopped at row one. The framework instead prices the option, tests it against the company's risk appetite, and hands development a precise target.

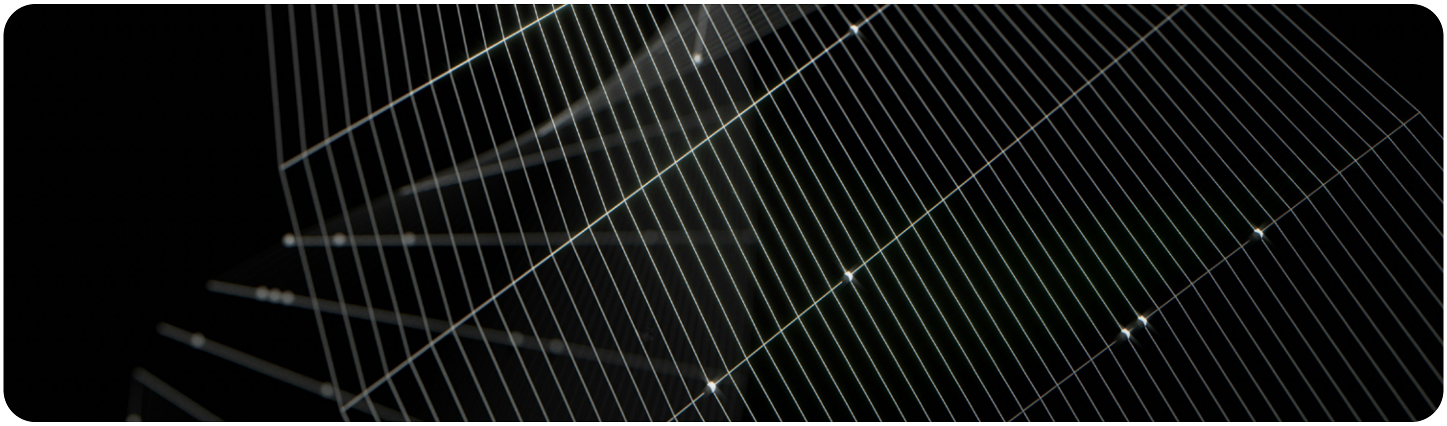

SCOPE AND LIMITS

What the method does, and doesn't, claim

The framework is a disciplined lens, not an oracle. Four caveats keep it honest.

It defines conditions, not verdicts.

The output is the set of conditions under which a project becomes valuable: the development targets and the budget that would justify them. It does not pronounce go or no-go. If the cheapest path to investability is implausible, for example, halving opex on a half-million-dollar budget, the model still reports it; whether that path is reachable is a judgment only the experts can make.

The option value is a transparent approximation.

The Real Options Model uses the Black-Scholes formula, which assumes a single future decision point and a smoothly distributed payoff. Real deployments can be staged and retimed, so the figure is best read as a defensible order of magnitude for what optionality is worth.

The reward test prices total risk on purpose.

The binding hurdle is a Sharpe-ratio test on the project's total volatility. That is a deliberate choice: An innovation is usually held as a concentrated bet the firm cannot diversify, so the risk that matters is all of it, not just the market-correlated part.

The outputs are only as sharp as the inputs.

Risk bands, market-correlation weights, and lever difficulties are expert estimates, and they drive the results. The framework makes these assumptions explicit, traceable, and consistent across the five stages, but it does not remove the judgment, and it does not make that judgment neutral. The experts closest to a project usually have an interest in seeing it funded, and optimism creeps in exactly where it matters most: Wider uncertainty inflates the option value, and softer difficulty weights make the targets look cheaper to hit. Precise-looking outputs still rest on these estimates.

This dependence on expert judgment makes independent data valuable.

The framework is most trustworthy when its inputs come from outside the room that wants the answer: objective technoeconomic benchmarks, peer cost and performance data, and learning-curve evidence drawn from the wider market rather than from a project's own advocates. Sourcing those inputs independently is precisely the role Lux Research is designed to play. The framework supplies the analytical rigor; independent, verifiable data ground that rigor in objective evidence.

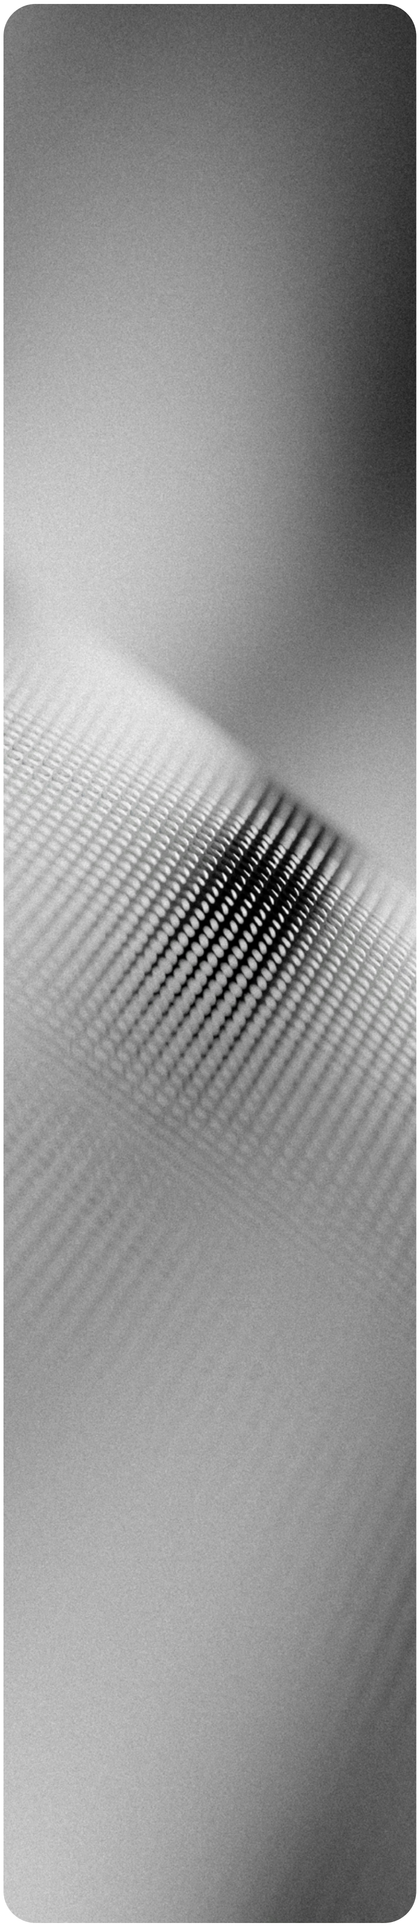

How to put the framework to use

The framework is most powerful when used as a single, connected workflow rather than five separate analyses. A handful of principles make the difference between running the models and changing the decision.

1. *Start before the case looks good.*

   The framework earns its keep precisely on innovations a deterministic gate would reject. A weak base case is the entry point, not the exit: the signal to quantify the upside rather than close the file.

2. *Value the option, not just the business case.*

   An innovation buys the right to invest later under truncated downside. Judge it by the value of that option and the development budget it justifies, not by an expected NPV that ignores the choice to walk away.

3. *Separate the development decision from the deployment decision.*

   The Real Options Model determines whether development is worth funding; the Investment Hurdles Model determines whether the eventual deployment justifies the capital. Keeping the two decisions separate prevents a marginal deployment from killing a valuable development option, and vice versa.

4. *Let your own risk appetite set the hurdle.*

   The return an innovation must earn is not universal. It depends on the funding company's beta, capex, and exposure. Anchor the hurdle to the firm that will ultimately deploy the technology, and resilience becomes a quantifiable advantage.

5. *Turn the analysis into targets.*

   The endpoint is not a verdict but an assignment: what to improve, by how much, and the budget that justifies it. Give the innovation team a brief grounded in the economics, and development effort focuses on the changes that alter the investment decision.

6. *Treat it as a living framework.*

   Costs fall, markets evolve, and uncertainty narrows as a program matures. Rerunning the workflow at each stage gate keeps the budget, the hurdle, and the development targets aligned with the project's evolving economics.

7. *Read the results across the portfolio.*

   This is a project-level method, but its outputs are designed to be compared. Apply the same workflow across candidate projects, and each produces a net option value, a justified development budget, and two risk-adjusted margins. Those outputs provide the basis for ranking projects and allocating a finite development budget. The comparison is only meaningful if the inputs are defined consistently: the same risk-band discipline, the same hurdle, and the same R&D-difficulty conventions. That consistency turns five per-project analyses into a portfolio decision — and is one of the strongest reasons to source the inputs independently rather than project by project.

Lux Take

Innovation is the deliberate purchase of an option under uncertainty.

The companies that win treat it that way: *funding the right to build, pricing risk honestly, and developing toward a target rather than a hope.*

The Lux Innovation Value Framework gives R&D and innovation leaders a defensible, end-to-end method for making the decisions that a deterministic business case cannot support. It does not promise that every innovation is worth funding. It promises that worthwhile innovations are no longer discarded by a model built for the wrong question and that those chosen for development come with a justified budget, a risk-adjusted hurdle, and clear technical targets.

Used together, the five stages of the framework transform uncertainty from the reason to reject an innovation into the source of its value.

About Lux

Lux Research fuels innovators to not only imagine what's possible in the future but also operationalize innovation success in the near term. We deliver research and advisory services to inspire, illuminate, and ignite innovative thinking that reshapes and grows businesses. Using quality data derived from primary research, fact-based analysis, and opinions that challenge traditional thinking, our experts focus on finding truly disruptive innovations that are also realistic and make good business sense.

The "Lux Take" is trusted by innovation leaders around the world, many of whom seek our advice directly before placing a bet on a startup or partner — our clients rely on Lux insights to make decisions that generate fantastic business outcomes. We pride ourselves on taking a rigorous, scientific approach to avoid the hype and generate unique perspectives and insights that innovation leaders can't live without.

For More Information

VISIT
www.luxresearchinc.com

READ
www.luxresearchinc.com/blog/

CONNECT
www.linkedin.com/company/lux-research

FOLLOW
@LuxResearch

CONTACT
questions@luxresearchinc.com

Lux Research | INSPIRE ILLUMINATE IGNITE