The Purse & the Knowledge

Why enterprise innovation allocation is an internal capital markets problem — and the questions a corporate partner must resolve before building the engine.

Variant A — The Internal Engine: enterprises allocating across innovation modes *without* a dedicated corporate venture arm.

Christian Gammill
Managing General Partner, Nextwave X Partners

A Framing Paper for Corporate Innovation Leadership
May 2026 · Version 6.0

THE ARGUMENT IN BRIEF

The instinct to run enterprise innovation like an internal venture fund is seductive and wrong. A fund works because it has a central node — one hub that holds the capital, the authority, and a clear view of every position. Strip that node away, as the modern enterprise does by default, and the metaphor collapses into a sprawl of distributed decisions, dispersed knowledge, and coupled bets that no one fully sees. Yet a center does exist: it lives on the CFO's line, where the power of the purse is real even though the knowledge is not. That precise shape — a center that controls the money but not the truth — is not novel. It is the ***internal capital market***, a structure corporate finance has studied for thirty years. Seen this way, the work also resolves into a prior clarity about ***what kind*** of capital each bet deploys, and an honest map of where discipline tends to fail. Reframing the problem this way changes what we are building. The task is no longer to design a better scorecard. It is to design a market — and to answer, deliberately, the institutional questions that decide whether that market creates value or quietly destroys it.

I The Seductive Wrong Metaphor

WHEN A LARGE COMPANY resolves to get serious about innovation spend, it almost always reaches for the same analogy: run it like a venture fund. Stage the bets. Follow on with the winners. Kill the losers early. Manage a portfolio rather than a wish list. The appeal is obvious — venture capital has spent forty years building exactly the discipline that enterprise innovation tends to lack, and the vocabulary travels well.

But the metaphor smuggles in an assumption that rarely survives contact with the enterprise. A venture fund is a **star topology**. At its center sits a single hub — the general partnership, expressed through an investment committee — that holds the capital, holds the authority to deploy it, and holds a clean, consolidated view of every position in the book. The spokes, the portfolio companies, are structurally independent: they do not share a balance sheet, a workforce, or a roadmap. One company's failure does not strand another's engineers. The hub decides; the spokes execute; the accounting is clean.

The discipline everyone admires in venture is downstream of that topology. It is the central node that makes ranking possible, that makes a kill enforceable, that makes a portfolio legible. Remove the node and you do not get a slightly messier fund. You get a different problem altogether.

II The Sprawl

The enterprise, by default, has none of the fund's structural gifts. Decision rights over innovation are distributed across business units, each with its own profit-and-loss statement and its own agenda. The binding resource is not fungible capital but lumpy, specialized people — the same engineers and product leaders who also keep the core business running, and who cannot be cleanly divided across three initiatives at once. The initiatives themselves are coupled: they share platforms and data, and one team's output is another team's dependency, so a bet cannot be valued in isolation. And no single node sees the whole board, let alone commands it.

Beneath the org-chart mess sits a deeper and older problem, and it is epistemic. The information that would make allocation rational — which initiative is actually working, where slack capacity is genuinely hiding, what a project truly costs once you price the distraction it imposes — is **local, tacit, and dispersed** across the people doing the work. It cannot be transmitted losslessly to a center. This is Hayek's knowledge problem, and it is the reason centralized innovation scorecards so reliably disappoint: they assume the relevant knowledge can be summoned to headquarters on demand. It cannot. Worse, the very act of summoning it distorts it — when the people who hold the knowledge know that funding hinges on what they report, what they report stops being knowledge and starts being strategy.

> ***A fund's discipline is downstream of its topology. The enterprise has the opposite topology — and reaching for the fund's tools without its structure is the original error.***

III The Center Reasserts Itself

And yet a center does exist — and naming it correctly is the turn on which everything else depends. In the structure that prompted this paper, the innovation-allocation function reports jointly to a chief innovation officer and to the chief financial officer. That second reporting line is decisive. To report to the CFO is to be plugged into the firm's actual capital-allocation machinery; it confers the one thing a pure convener lacks, which is the **power of the purse**.

So the topology is neither of the two we have considered. It is not a venture fund — there is no hub with both full knowledge and full authority. But neither is it a pure commons with no center at all. It is the precise, well-mapped shape in between: **a center that holds the purse but not the knowledge, allocating to units that hold the knowledge but not the purse.**

Corporate finance has a name for this structure and three decades of study behind it. It is the **internal capital market** — the mechanism by which a diversified firm moves scarce capital across its divisions, and by which headquarters earns (or forfeits) its keep as an allocator. The intellectual home for the enterprise-innovation problem is not the venture literature at all. It is the work of Stein, of Scharfstein and Stein, of Rajan, Servaes and Zingales on how multi-business firms allocate internal funds, and why they so often allocate them badly.

The reframe carries an immediate consequence for the person at the center. The head of innovation is *not* a general partner. They are the operator of an internal capital market — and the right mental model for that role is closer to a central banker or a market designer than to a fund manager. Their leverage is not in picking every bet. It is in setting the currency, the rules, and the visibility within which distributed actors make and unmake bets.

IV What Kind of Capital? The Deployment Continuum

So far we have spoken of "capital" as if it were a single substance. It is not — and the first discipline an internal market imposes is to name *what kind* of capital a given bet actually deploys. Innovation capital runs along a continuum, from cheap, reversible optionality at one end to expensive, irreversible control at the other.

LOW COMMITMENT · OPTIONALITY | HIGH COMMITMENT · CONTROL

01	02	03	04	05
Explore	**Build**	**Partner**	**Invest**	**Buy**
Sensing & optionality	Owned capability	Shared capability	Financial position	Acquired capability

REVERSIBILITY FALLS · INFLUENCE COST OF A KILL RISES →

The five modes are not mutually exclusive; most enterprises run a portfolio across them. Read through the lens of this paper, the continuum is also a map of where allocation discipline will fail.

The categories are familiar enough, and most enterprises run several at once. The point is not the taxonomy but what it reveals when laid against the argument so far. **The continuum is a gradient of reversibility, and reversibility governs how hard a bet is to kill.** An Explore-stage probe is cheap and optional; unwinding it costs little, and almost no one's standing depends on it. A full acquisition or a multi-year owned build is the opposite — capital is sunk, careers are attached, and the commitment is, for practical purposes, irreversible.

This produces an uncomfortable alignment. *Influence cost rises in lockstep with commitment.* The bets that are hardest to kill are precisely the ones at the high-commitment end — exactly where the most capital is concentrated, and exactly where a bad position does the most damage. The continuum therefore predicts where the socialism of an internal capital market will bite hardest. It is not merely a menu of instruments; it is the risk map of the allocation discipline that follows.

V The Dark Side: Why the One Move That Matters Most Is the One These Markets Fail At

The promise of an active internal capital market is what the literature calls **winner-picking**: headquarters adds value by moving resources from weak claims to strong ones faster, and with better information, than an external market could. That is the bull case, and it is real. An active internal market beats a passive one precisely because the center can reallocate.

The empirical record, though, is sobering — and it should govern the entire design. Internal capital markets exhibit what corporate finance bluntly names **socialism**: capital gets spread too evenly, strong divisions are taxed to subsidize weak ones, and allocation tracks divisional *power* rather than marginal return. The mechanism is influence. Divisions expend real effort lobbying for resources and defending their existing claims, and that rent-seeking — the cost of influence activity — systematically bends the allocation away from where value would say it should go.

The implication for the tool is sharp, and it is counterintuitive. Your two headline moves are not symmetric. The **double-down** is the easy half: everyone enjoys backing a winner, and a rising bet attracts capital almost on its own. The **kill** is the hard half — and it is, specifically and predictably, the move internal capital markets are worst at. Killing is where influence costs concentrate. It is where the threatened sponsor fights hardest, where sunk cost and personal reputation entangle, where the politics turn ugliest. A portfolio's quality is set far more by its willingness to withdraw resources from losers than by its enthusiasm for winners — and that willingness is exactly the thing the structure conspires to erode — and, as the continuum makes plain, the erosion is worst exactly where the stakes are highest.

> ***The double-down is the easy half. The kill is where influence costs concentrate — which is why it is the design problem, not a feature to bolt on.***

This relocates the center of gravity of the whole build. The hard part is not the scoring model that ranks initiatives; ranking is comparatively easy. The hard part is the machinery that lets the center withdraw capital and people from a losing claim *against the resistance of the claimant*. Three mechanisms do most of that work.

Capital that is re-earned, not entitled

The default state of resources should be recycling, not renewal. If a budget, once granted, is treated as an annual entitlement that must be actively pried away, influence wins by inertia. If capital must instead be re-earned each cycle against current evidence, the burden of proof sits where it belongs — on the claim, not on the allocator.

Kill criteria pre-committed at funding

The thresholds that would end an initiative should be written down at the moment of funding, before any ego is attached to the outcome. A kill that fires as a pre-agreed rule is survivable; a kill imposed later as a fresh judgment invites the full machinery of influence. Pre-commitment converts a political act into the execution of a contract.

A revelation layer that scores the forecaster, not just the forecast

Because funding depends on projections, projections will be shaded — this is structural, not a matter of bad actors. The countermeasure is to track each sponsor's *calibration* over time and price it in, so that chronic optimism becomes visibly expensive to its source. In effect, the system runs a Bayesian update on people, not only on projects: a sponsor whose forecasts have repeatedly missed should see their next forecast discounted automatically. This is the single most neglected component of enterprise innovation systems, and the one that most directly attacks the corruption of information at the source.

VI Two Masters, Two Currencies

The dual reporting line is not an accident of the org chart to be tidied away. It is the **explore/exploit tension** — the oldest problem in the study of organizational learning — wired directly into governance. The CFO wants discipline, return, and capital efficiency. The chief innovation officer wants optionality, learning, and the chance at transformation. These are different objective functions, and a serious tool must serve both without letting one quietly consume the other.

A single blended expected-value score cannot do this, and the failure is not cosmetic. Collapsing everything to one number forces a transformational bet to compete against a cost-savings project on the one axis the moonshot is guaranteed to lose: near-term, quantifiable financial return. The predictable result is a portfolio that believes itself balanced while systematically starving its own future — every cycle, the safe bet out-scores the consequential one, and the firm optimizes itself into irrelevance one defensible decision at a time.

The resolution is **two currencies, separately budgeted**. Near-term bets are judged in the CFO's coin — risk-adjusted financial return, held to real discipline. Transformational bets are judged in a different coin entirely: option value and validated learning. Did this bet cheaply resolve a large uncertainty? Did it open a door that was previously closed? The two pools are sized deliberately and do not compete inside a single ranking. Crucially, this makes ***balance*** a conscious structural decision about the relative size of the pools — an explicit statement of strategy expressed in numbers — rather than an emergent and usually accidental property of one leaderboard.

The two currencies and the deployment continuum are orthogonal, and they fit together cleanly: the continuum names the ***instrument***, the currency names the *test*. Bets toward the Explore end are judged mostly in the coin of option value and learning; bets toward the Buy end in the coin of financial return; the middle blends the two. One axis says what kind of capital is at work; the other says how to know whether it is working.

VII The Central Questions

There is a layer that sits above even these mechanism questions, and it has to be set first: the enterprise's *posture*. Before a single initiative is scored, leadership must state where it stands on five dimensions — over what horizon, led by whom, through which capital modes, toward which goals, and on what terrain — the five dimensions below. Skipping this step is the most common failure of all; in its absence, project scoring quietly optimizes against an allocation logic no one has actually agreed to.

TIME	WHO	HOW	WHY	WHERE
Commitment horizon	**Leadership mix**	**Capital mode**	**Program goals**	**Focus areas**
0–1 yr · 2–4 yr · 5+ yr	Inside vs outside · employee vs advisor	Explore · Build · Partner · Invest · Buy	Ecosystem · revenue · venture returns	Core vs adjacent · industry vs competency

The strategic setup: five dimensions resolved before any initiative is scored. Set them first and every downstream choice inherits a stated logic; leave them implicit and the scoring optimizes against an allocation posture no one has agreed to.

With posture set, the remaining work is institutional design, and it turns on a small number of decisions the corporate partner must make consciously and early. None of these is a technical question about the tool; each is a question about authority, incentives, and truth. They are, in rough order of consequence:

1 **Is the capital actually mobile, or merely budgeted?**
The decisive question. Can the center genuinely reclaim and redeploy funds mid-stream, or are budgets annual entitlements that turn untouchable the moment they are granted? Active internal capital markets create value; passive ones earn the diversification discount. If capital cannot move, what is being built is a tracker, not an engine — and no amount of elegant scoring will change that.

2 **Are people governed by the same authority as the dollars?**
The scarce resource is specialized talent, and it is usually controlled by the business units rather than the allocator. If capital is mobile but engineers are not, the engine steers with one hand on the wheel. The partner must decide whether the innovation allocator holds any real claim on people, or only on budget — because a reallocation that cannot move the team is a reallocation in name only.

3 **How is truth revealed, and who bears the cost of being wrong?**
Absent a revelation mechanism, the inputs are adversarial by construction. Does a sponsor carry skin in the game? Is forecast accuracy tracked, remembered, and made consequential? Without this, the most optimistic projector wins the most capital — which is exactly backwards, and exactly how internal markets drift toward funding confidence instead of merit.

4 **What, precisely, is the center's authority — purse, people, or only convening?**
Most heads of innovation occupy an uncomfortable middle: real influence, partial budget, no claim on talent. The correct mechanism design is entirely different at each point on that spectrum — a command allocator over a carved-out pool is a different machine from an incentive market over resources one can only nudge. This must be named explicitly, not assumed, because the whole architecture follows from it.

5 **How are the two currencies sized, and who sets the exchange rate?**
Someone must decide what fraction of resources is governed by financial discipline versus optionality. That decision *is* the firm's innovation strategy, expressed in numbers — and it should be made deliberately and jointly by the CFO and the chief innovation officer, not left to drift out of a hundred uncoordinated funding calls.

6 **What makes a kill legitimate?**

Because killing is the hard part, the partner must settle in advance what authority, what evidence, and what ritual make a withdrawal of resources actually stick against a determined sponsor. A kill switch with no legitimacy is a button no one will ever press — and an innovation portfolio that cannot kill is simply a slower way of spreading capital too thin.

VIII From Scorecard to Market

The value of the reframe is that it relocates the problem. Treated as a measurement exercise, enterprise innovation asks a thin question: how do we rank our initiatives? Treated as an internal capital market, it asks a far richer and more honest one: **what rules allow a center that holds the purse, but not the knowledge, to allocate well — against agents who know more than it does and want more than it can give?**

That second question is the one worth building a tool around, because it refuses the comfortable fiction that better data alone will fix allocation. It will not. The dispersion of knowledge is permanent, the influence pressure is structural, and the dual mandate is real. What can be designed is the institution that metabolizes all three: a market with a mobile currency, a truthful revelation layer, a legitimate kill, and an honest accounting of optionality alongside return. Build that, and the scorecard becomes a detail. Skip it, and the most sophisticated scorecard in the world will still fund the loudest division.

The operating layer that answers all of this has a recognizable shape. It is not a project-scoring tool; it is a **decision system** — opinionated about allocation, structured about evidence, and auditable about why. It sets posture before it scores anything; it places every bet on the continuum rather than in an undifferentiated pile; it holds each initiative's rationale and evidence in a durable record instead of a forgotten deck; and it lets outcomes flow back to sharpen the next allocation, not merely settle this one. A system of exactly that shape is the premise behind *Decided.ai* — and the argument of this paper is the case for why an enterprise needs one.

This is the terrain Nextwave X Partners works in by conviction — the conviction that allocation under deep uncertainty is won or lost on institutional design, not on the precision of any single estimate. The questions above are not preliminaries to the build. They *are* the build. Settle them deliberately, and the engine follows.

SELECTED REFERENCES — INTELLECTUAL LINEAGE

1. Hayek, F. A. "The Use of Knowledge in Society." *American Economic Review* 35, no. 4 (1945): 519–530.
2. Stein, Jeremy C. "Internal Capital Markets and the Competition for Corporate Resources." *Journal of Finance* 52, no. 1 (1997): 111–133.

3. Scharfstein, David S., and Jeremy C. Stein. "The Dark Side of Internal Capital Markets: Divisional Rent-Seeking and Inefficient Investment." *Journal of Finance* 55, no. 6 (2000): 2537–2564.
4. Rajan, Raghuram, Henri Servaes, and Luigi Zingales. "The Cost of Diversity: The Diversification Discount and Inefficient Investment." *Journal of Finance* 55, no. 1 (2000): 35–80.
5. Milgrom, Paul, and John Roberts. "An Economic Approach to Influence Activities in Organizations." *American Journal of Sociology* 94 (1988): S154–S179.
6. March, James G. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2, no. 1 (1991): 71–87.
7. O'Reilly, Charles A., and Michael L. Tushman. "The Ambidextrous Organization." *Harvard Business Review* (April 2004).
8. Ostrom, Elinor. *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge University Press, 1990.

ABOUT THE AUTHOR

Christian Gammill has worked at nearly every point on the continuum this paper describes. He was raised inside the most structured innovation engine there is — pharmaceutical drug discovery — as the son of a PhD medicinal chemist and a pharmaceutical IP attorney, and trained as a chemist himself, doing method development at a contract research organization serving big pharma. At IBM he helped build a data model to define and activate innovation across a global enterprise, deconstructing it from notional concepts down to measurable, actionable drivers through means-end chain analysis and structural equation modeling, in service of a Bayesian-inspired decision framework. He then spent years in early-stage technology startups, and today, as Managing General Partner of Nextwave X Partners, helps large enterprises deploy venture capital into innovation. The failure modes described in these pages are observed, not theorized.