


BEYOND THE APPROVAL GATE

Governing What You Don't Yet Understand

Why the machinery built to control known work quietly strangles innovation — and the different discipline governance must learn



Giovanni Leonardi

May 2008

Contents

1. Executive Summary	3
2. The committee that asked the wrong question	3
3. Why our instruments mislead us	4
4. The two economies inside one organisation	5
5. Governing uncertainty without abandoning it	6
6. The objection worth taking seriously	7
7. The discipline the committee has to learn	8

Governing What You Don't Yet Understand

Transformation · No. 175

Originally written May 2008 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The instrument survives not only because it looks rigorous but because it launders individual risk.”

1. Executive Summary

Every organisation of any ambition now keeps two kinds of work under one roof: the work it understands and the work it is trying to understand. The first kind — the upgrade, the rollout, the consolidation — can be forecast, costed, and held to a plan. The second kind — the genuinely new product, channel, or capability — cannot, because its defining feature is that nobody yet knows how it will behave. The difficulty this essay examines is that most organisations govern both kinds with a single instrument, built for the first, and are then surprised when it quietly kills the second.

The machinery of oversight we have inherited — the business case with its three-year payback, the capital-approval committee, the phase-gate that demands certainty before it releases money — is an achievement, not a mistake. It was built to stop good money chasing bad, and it does. But it selects for forecastability, and forecastability is inversely related to novelty. Pointed at innovation, it does not neutrally assess; it systematically funds the least innovative options and starves the rest, while producing a paper trail that makes the whole process look rigorous.

The argument here is not that innovation should escape governance. Ungoverned innovation is not brave, it is merely expensive, and the practitioner who has watched an "innovation fund" evaporate into a dozen unaccountable pet projects knows this as well as any auditor. The argument is that governing the unknown is a different discipline from governing the known, with different questions, different evidence, and a different definition of what a good decision looks like. It commits in stages rather than in advance; it buys the right to continue rather than the obligation to deliver; it treats a well-run failure as information purchased at a fair price rather than as a control breakdown. Holding both disciplines inside one organisation, and being honest about which one a given piece of work belongs to, is the real task. This essay sets out why the single-instrument habit persists, what it costs, and what the second discipline actually requires of the people who sit on the committee.

2. The committee that asked the wrong question

Picture the scene, because most of us have sat in some version of it. A capital-approval committee, quarterly, competent, conscientious. On the agenda, among the server refreshes and the ERP extension, sits a proposal for something the organisation has never done before — call it a new self-service channel for a customer segment the firm has never served directly. The sponsor is asked, reasonably enough, for the same thing every other item on the agenda has provided: a business case with a three-year cash flow, a payback period, and a sensitivity analysis.

The sponsor does the honest thing, which is the fatal thing. She admits that she does not know the take-up rate, because no one has offered this segment anything like it before. She does not know the unit economics at scale, because there is no scale yet. Her numbers, she says, are assumptions dressed as forecasts. The committee, being diligent, marks the proposal down for exactly the quality it should have marked it up for: candour about its own uncertainty. The server refresh, whose numbers are solid because the work is understood, sails through. The channel is sent away to "firm up the business case" — which is to say, to manufacture a false confidence it does not possess — and by the next quarter the sponsor has learned the lesson the system teaches. She comes back with crisp, defensible, entirely fictional numbers, and the committee, satisfied, approves a plan that everyone in the room half-knows is a work of fiction.

Consider a composite that is truer to life than any single case. An organisation stands up an innovation fund — a genuine act of intent — and endows it with eight million over three years. It governs the fund through the existing capital process, because that is the process the finance function trusts. Three years on, the money is spent and the post-mortem is revealing. Of the fourteen initiatives funded, eleven were shelved before launch. Not one of the eleven was stopped because it had been tried and had failed in the market; every one was stopped earlier, at a gate, for failing to produce a business case robust enough to survive the next tranche of funding. The three that survived to launch were, on inspection, the three least novel — extensions of things the firm already did, whose returns could be forecast precisely because they broke no new ground. Roughly six of the eight million had gone to work that the ordinary budget would have funded anyway. The fund had not bought innovation. It had bought a more elaborate way of doing the familiar, wrapped in the language of the new.

The gate did exactly what it was designed to do: it selected for certainty. The tragedy is that in the domain of innovation, certainty and value point in opposite directions. The proposals that could prove their returns were the ones least worth funding; the ones worth funding could not, by their nature, prove anything yet.

The committee did nothing wrong by its own lights. That is the uncomfortable part. Every individual decision was defensible. The failure was systemic and invisible, located not in any single judgement but in the instrument itself.

3. Why our instruments mislead us

To see why this happens, it helps to be precise about what a conventional business case actually is. It is a machine for converting a description of the future into a decision in the present, and it runs on a specific fuel: reliable estimates. Feed it good estimates and it produces good decisions with real accountability — the sponsor commits to a number, and later we can check whether the number was met. This is not bureaucracy for its own sake; it is one of the genuine achievements of modern management, and in the wake of the governance scandals earlier this decade, boards have every reason to prize the discipline it imposes.

But notice the hidden assumption. The machine treats an estimate as a forecast — a claim about a knowable future that can later be judged right or wrong. For understood work, that assumption holds. The migration will take four months or it will not; if it takes seven, someone estimated badly and we can learn from it. For genuinely novel work, the assumption is simply false. The take-up rate of a channel that has never existed is not a

forecast that turns out right or wrong; it is a guess about a system that does not yet exist and whose behaviour will only be revealed by building it. Demanding a forecast where only a guess is possible does not produce better information. It produces the same guess, laundered through a spreadsheet until it acquires the appearance of fact.

The mechanism, then, runs in three steps, and every practitioner should be able to recite it:

1. The gate demands forecastable returns as the price of funding.
2. Forecastability falls as novelty rises — the newer the thing, the less honestly it can be forecast.
3. Therefore the gate, applied evenly, funds in inverse proportion to novelty, while appearing to be neutral.

The appearance of neutrality is what makes this so durable. Nobody stands up in the committee and says "let us avoid anything genuinely new." On the contrary, everyone sincerely wants innovation. But the instrument does the discriminating quietly, on their behalf, dressed as prudence. And because the instrument produces a clean audit trail — every rejection documented, every approval justified — the whole apparatus is not merely tolerated but rewarded. It looks like exactly the rigorous, defensible governance that a well-run board is supposed to want. This is why the pattern is so stubborn across organisations and across the decade: it is not a lapse in control but an excess of the wrong kind of control, and excess control is very hard to argue against in a boardroom that still remembers what too little control cost.

There is a second, subtler force at work. The people who staff these committees are, rightly, held accountable for the money they release. When a forecast-backed project fails, the sponsor is answerable, but the committee can point to the process it followed. When a proposal is approved on the frank admission that its returns are unknowable, and it then fails, the committee feels exposed in a way the process does not protect against. The rational individual response is to demand forecasts even where forecasts are meaningless, because a forecast transfers the blame. The instrument survives not only because it looks rigorous but because it launders individual risk. We should not moralise about this. It is a predictable response to how accountability is currently drawn, and any serious remedy has to change the accountability, not merely exhort people to be braver.

4. The two economies inside one organisation

The way out begins with an admission that many organisations resist: that they are running two different economies under one set of books, and that these economies obey different laws.

The first is the economy of exploitation — extracting value from what the organisation already knows how to do. Here, efficiency is the virtue. Variance is waste. A good outcome is the plan, delivered. The business case, the gate, and the variance report are precisely the right instruments, because the work is knowable and the point of governance is to hold it to what was known.

The second is the economy of exploration — searching for what the organisation does not yet know how to do. Here the virtues invert. Variance is not waste; it is the entire point, because you are buying information you did not previously have. A good outcome is not "the plan, delivered" — there was no reliable plan — but "the cheapest possible discovery of whether this works." In exploration, a project that runs exactly to its original assumptions has usually learned nothing, which is its own kind of failure.

The scholarship of the last few years has given this pairing a name — the ambidextrous organisation, one that can exploit and explore at the same time — and the observation underneath it is old practitioner wisdom: the disciplines that make you excellent at running the known business are the very disciplines that make you hopeless at building the next one. The error most organisations make is not that they lack one economy or the other. It is that they govern both with the instruments of the first. They take a control system tuned for variance-suppression and point it at an activity whose whole purpose is to generate variance, and then they wonder why nothing new survives contact with the committee.

“Variance, in the economy of exploration, is not a defect to be suppressed. It is the product you are paying for.”

Naming the two economies is not a semantic nicety. It is the precondition for everything that follows, because it lets us ask, of any given piece of work, the only question that matters at the outset: which economy does this belong to? Almost all the damage I have seen comes from misclassification — from work that is genuinely exploratory being waved through the exploitation machinery because no one thought to ask which kind of thing it was. Get the classification right and the governance almost designs itself. Get it wrong and no amount of committee diligence will save you.

5. Governing uncertainty without abandoning it

If the conventional instrument is the wrong tool for exploration, what is the right one? Not the absence of governance — that merely trades one failure for a worse one — but a governance calibrated to uncertainty rather than offended by it. Its logic differs at almost every point from the capital-approval model, and the differences are worth setting out plainly.

Governing the known	Governing the unknown
Approve the whole business case up front	Fund the next increment of learning only
Evidence is a credible forecast	Evidence is what the last increment revealed
Success is the plan, delivered	Success is uncertainty, cheaply reduced
A stopped project is a failure of delivery	A stopped project is a completed experiment
The gate asks: will this hit its numbers?	The gate asks: have we learned enough to justify more?
Variance is a control problem	Variance is the information you paid for

The organising idea beneath the right-hand column is one that finance functions already understand in another guise: the idea of an option. When the future is uncertain, you do not buy the whole outcome in advance; you buy the right, but not the obligation, to continue — and you buy it in the smallest tranche that will tell you something. Each round of funding purchases not a deliverable but a reduction in uncertainty. The governing body's job is no longer to judge a forecast it cannot honestly evaluate, but to judge something it can: did the last increment of spending buy a proportionate increment of

knowledge? Is the remaining uncertainty now small enough, and the potential prize large enough, to justify the next tranche?

This reframing changes what the committee asks for, and it is worth being concrete about the shape of it, because vagueness here is how these schemes collapse back into the old habits:

- Fund learning, not deliverables. The unit of investment is the experiment that resolves the biggest current unknown for the least money — not the finished product. The first tranche for our new channel does not build the channel; it buys the answer to the one question whose answer would kill the idea fastest if it were bad.
- Set the gate on knowledge, not on plan-adherence. At each review the question is not "are you on track against the business case?" — there is no reliable business case — but "what did we just learn, and does it justify continuing?" A team that has learned its idea will not work, and says so, has succeeded at that gate, not failed.
- Cap the downside deliberately. Uncertainty is survivable when exposure is bounded. The discipline is to make each tranche small enough that being wrong is affordable, so the organisation can afford to be wrong often, which is the only way it will occasionally be right about something large.
- Pre-agree what would kill it. The most useful thing a sponsor can bring is not a forecast of success but a clear statement of the evidence that would make them stop. Naming the kill-criteria in advance is what separates disciplined exploration from a pet project that never dies.

None of this is soft. If anything it is harder than the old model, because it removes the sponsor's hiding place. Under the forecast model, a failing project can limp along for years, always "on plan" against numbers everyone privately disbelieves. Under staged, option-based funding with pre-agreed kill-criteria, a project that is not learning is stopped early and visibly. Properly run, this discipline kills more projects than the conventional one, and kills them sooner and cheaper. That is not a weakness of the approach. It is the whole point. The organisation that cannot bring itself to stop things has not escaped the failure of the innovation fund; it has merely relocated it.

6. The objection worth taking seriously

An essay that only argued with a straw man would not be worth the reading, and the objection to all this is a strong one, so let me put it at its strongest rather than at its weakest.

The objection runs: this is a licence for unaccountable spending. Strip away the language of options and learning and what remains, a hard-nosed finance director might say, is a request to hand money to people who cannot say what they will deliver, cannot forecast a return, and now — conveniently — have a theory that reclassifies their failures as "learning." The forecast-and-gate model may be blunt, but at least it is honest about accountability: someone commits to a number and is answerable for it. Loosen that and you do not get more innovation; you get the same drift and waste that discredited "innovation funds" in the first place, only now with an intellectual justification that makes it harder to challenge. This is a serious argument, and anyone who has watched money disappear into a fog of "strategic learning" with nothing to show for it should feel its force.

It deserves a serious answer, not a dismissal, and the answer is this: the option-based model is not less accountable than the forecast model. It is accountable to a different and more honest measure. The forecast model holds the sponsor to a number that both parties

know to be fictional — an accountability that is rigorous in form and empty in substance, because everyone can see the number was invented to pass the gate. The option model holds the sponsor to things that are real and checkable: Did you spend only the tranche you were given? Did you run the experiment you said you would run? Did it resolve the uncertainty it was meant to resolve? Are you honouring the kill-criteria you set, or moving the goalposts to keep the project alive? These are not soft questions. They are, if anything, harder to wriggle out of than a three-year forecast, because they concern what actually happened last quarter rather than what is promised for a quarter three years away that no one will remember to check.

The distinction that resolves the objection is between unaccountable spending and uncertain spending. The finance director is entirely right to refuse the first and entirely wrong to conflate it with the second. Unaccountable spending has no bounded exposure, no defined experiment, no evidence test, and no pre-agreed point at which it stops. Uncertain spending has all four; what it lacks is a reliable forecast — and demanding a reliable forecast of a genuinely uncertain thing is not accountability, it is theatre. The whole craft of governing innovation lies in supplying the four disciplines the objection rightly demands while refusing the one demand the objection wrongly insists on. A firm that cannot tell the difference will oscillate between two failures: strangling innovation in the name of control, or funding drift in the name of boldness. The point of the second discipline is to have neither.

7. The discipline the committee has to learn

If there is a single thing to carry away from all this, it is that governing innovation is not a matter of relaxing governance but of changing the question the governance asks. The conventional committee asks, of everything that comes before it, "will this deliver the returns it promises?" — and for the known business that is exactly the right question. For the unknown business it is the wrong one, and asking it faithfully is precisely how the unknown business gets killed. The right question for exploratory work is narrower and more honest: "have we learned enough, at a cost we can bear, to justify learning a little more?"

Almost everything else follows from getting that question right. The classification of work into the two economies follows from it, because you cannot ask the right question until you know which kind of work you are looking at. The staged, bounded, option-shaped funding follows from it, because that is simply what it looks like to buy learning rather than to buy outcomes. And the different, sharper form of accountability follows from it too, because "what did you learn for the money?" turns out to be a more searching question than "did you hit your forecast?"

The organisations that manage this do not have braver committees or looser controls. They have committees that have learned to hold two disciplines at once, and — this is the hard part — to be honest, in each case, about which discipline applies. That honesty is the whole of the craft. It is easy to run the known economy well and easy, in a different way, to indulge the unknown one carelessly. What is difficult, and rare, and worth building, is the organisation that can do both under one roof without letting the instruments of the one quietly destroy the other. The committee that learns to ask which economy it is in, before it asks anything else, has already done the larger part of the work.

We are, most of us, fluent in the governance of the known. We have spent a decade becoming more so, and for good reason. The task now is to become even half as fluent in the governance of the unknown — and to stop mistaking our fluency in the first for

competence in the second.