

BEYOND AUTOMATED ALLOCATION

The Portfolio That Rebalances Itself — and Why You Shouldn't Let It

The Case for Automating Insight but Never Allocation in Portfolio Governance

Giovanni Leonardi

March 2026



The Portfolio That Rebalances Itself — and Why You Shouldn't Let It

Portfolio · No. 122
March 2026

“The moment the system starts making the decision as well as informing it, the prize becomes a trap.”

The Seductive Promise of the Closed Loop

The quarterly portfolio review was still three weeks away when the platform flagged it: two initiatives running fourteen percent over forecast burn, a third whose projected benefits had decayed below threshold, and a recommended reallocation that would redistribute £4.2 million across four programmes by the end of the month. The analysis was defensible. The numbers were clean. The recommendation was wrong.

It was wrong not because the model had miscalculated — it had not — but because the decision it proposed to automate was never a calculation in the first place.

We are now deep into the era of AI-assisted portfolio management. The tooling has matured rapidly over the past eighteen months, and the capabilities are genuine: continuous ingestion of delivery signals, real-time benefits modelling, dependency mapping that updates as plans shift. What was once a quarterly exercise in spreadsheet archaeology has become a persistent, machine-driven surveillance of portfolio health. This is an unambiguous good. The danger lies not in the sensing but in the step that follows it.

The pattern I see recurring — across sectors, across maturity levels — is the same: organisations that begin with AI as an early-warning system quietly migrate towards AI as a decision-maker. The dashboard that began by surfacing risk graduates to proposing action, then to proposing action with a deadline, then to proposing action with a deadline and a pre-populated change request. Each step is individually reasonable. The trajectory is not.

What the Machine Reads Well

To understand why the closed loop fails, we first need to acknowledge — with some admiration — what it does extraordinarily well. AI-assisted portfolio tools are genuinely superior to human review at a specific class of signal:

- Delivery velocity and trajectory. Models trained on earned-value and throughput data detect slowdowns weeks before a programme manager's status report

acknowledges them. The signal is in the gradient, not the snapshot — and machines read gradients naturally.

- Benefits decay. When a business case assumed a market window, a regulatory deadline, or a competitive gap, the model can track the external conditions against the original assumptions and flag when projected benefits have eroded. This is work that humans defer endlessly because it requires revisiting uncomfortable truths.
- Resource contention patterns. Cross-portfolio resource clashes — the same enterprise architect committed to three critical paths, the same test environment needed by four teams in the same sprint — are invisible in any single programme's view and obvious in the portfolio-wide data.
- Dependency risk propagation. When Programme A's delay cascades into Programme B's integration milestone, which in turn threatens Programme C's go-live window, the chain is computable. Most organisations discover it in a crisis meeting. The model discovers it on a Tuesday morning.

These are real capabilities, and they represent a genuine advance over the quarterly review cycle that preceded them. The mistake is to assume that because the machine can sense these signals faster, it should also act on them faster.

What the Machine Structurally Cannot See

Portfolio decisions are not optimisation problems. They are governance acts — and governance encodes things that do not appear in delivery data.

Risk appetite is a judgement, not a metric. When an organisation decides to continue funding a programme that is over budget and behind schedule, it may be making a perfectly rational choice — because the strategic cost of abandoning that initiative outweighs the delivery cost of persisting with it. The AI sees the burn rate. It does not see the board conversation in which the chief executive committed to the regulator that this capability would be delivered. It does not see that withdrawing from this programme would signal to a joint-venture partner that the organisation is retreating from a market. Risk appetite is set in rooms the model has never entered, against considerations the data does not encode.

Strategic bets are deliberately inefficient. A well-governed portfolio contains initiatives that look suboptimal by any delivery metric because they are bets — investments in learning, in positioning, in optionality. The innovation programme that has consumed eighteen months and produced no deployable output may be exactly on track if its purpose is to build organisational capability in a technology the leadership believes will matter in three years. The model sees eighteen months of cost with no measurable benefit. The strategy sees a hedge against irrelevance.

Political commitments are load-bearing. We may wish this were not so, but portfolio decisions exist within a web of commitments — to ministers, to regulators, to unions, to partners — that constrain the decision space in ways no dataset captures. A reallocation that is optimal on paper may be impossible in practice because it would break a promise that holds a coalition together. This is not dysfunction; it is the reality of governing complex organisations.

Stability is itself a deliverable. Perhaps the most insidious effect of continuous rebalancing is what it does to the teams doing the work. Delivery requires commitment horizons — teams need to know, for some reasonable period, that their programme will not be defunded, their scope will not be redirected, their people will not be pulled. Continuous portfolio optimisation, even when each individual adjustment is small, creates an environment of permanent contingency. People stop investing in hard problems because they have learned that the portfolio may pivot before the solution matures. The machine optimises the allocation; the organisation loses the capacity to deliver anything that takes longer than one rebalancing cycle to complete.

The Design Principle

The right architecture is not difficult to state, though it requires discipline to hold: automate the sensing, never the committing.

Let the AI tools run continuously. Let them ingest every delivery signal, every benefits update, every resource contention pattern. Let them surface anomalies the moment they appear. Let them model scenarios — if we reallocate here, the projected impact is this; if we hold, the risk trajectory is that. This is where machine speed genuinely helps: the time between a signal emerging and a human being aware of it should approach zero.

But the decision to act on that signal — to reallocate funds, to pause a programme, to accelerate an initiative — must remain a human act, taken at a human rhythm, inside a governance structure that can weigh what the data shows against what the data cannot show. Not because humans are better at optimisation — we are manifestly not — but because portfolio decisions are not optimisation. They are acts of institutional judgement that integrate information the model does not possess and priorities the model cannot rank.

The goal is not faster decisions. The goal is faster awareness feeding decisions made at the tempo governance requires — machine-speed sensing, human-speed committing.

This means designing the human decision cadence deliberately. If the old quarterly review was too slow — and it was — the answer may be monthly, or six-weekly, or triggered by threshold breaches that the AI itself defines. But whatever the cadence, it must be a cadence: a rhythm that gives delivery teams stability, gives decision-makers time to consult beyond the data, and gives the organisation the predictability it needs to execute.

The portfolio director who walks into a monthly review armed with six weeks of continuous sensing, pre-modelled scenarios, and clearly flagged anomalies is in a profoundly better position than one working from last quarter's spreadsheet. That is the prize. The moment the system starts making the decision as well as informing it, the prize becomes a trap.

The Discipline of the Split

The objection, of course, is speed. In a world where conditions shift weekly, can we really afford to wait for the next scheduled review? The answer is that we cannot afford not to. The cost of a delayed reallocation — real as it is — is almost always smaller than the cost of an organisation that has lost its ability to commit. Rapid sensing closes most of the speed gap in any case: the portfolio board that meets monthly with continuous AI-driven intelligence is not slow. It is, by any historical standard, extraordinarily fast. What it is not is automatic — and that is precisely the point.

We are at a moment where the capability of the tooling is running ahead of our governance imagination. The platforms can do more than we should let them do, and the vendors, understandably, frame full automation as the destination rather than the boundary. The organisations that will govern their portfolios well in this environment will be those that draw the line clearly and hold it: automation of insight, never of allocation. It is not the most exciting position to take. But it is the one that preserves the thing a portfolio governance system must protect above all else — the capacity of the organisation to make decisions that are genuinely its own.