

BEYOND FAIR WEATHER

Crisis as Catalyst: Why Most Prioritisation Frameworks Are Fair-Weather Tools — and What to Build Instead

What the downturn's forced triage revealed about portfolio governance, and the discipline organisations should keep once the pressure lifts



Contents

1. Executive Summary	3
2. The prioritisation that only happened under a gun	3
3. Why the frameworks failed when they were finally needed	4
4. The evidence: what the cut actually revealed	4
5. What genuinely worked under constraint	5
6. The options: how to keep what the crisis taught	6
7. The objection worth taking seriously	6
8. Recommendation	7
9. What it would take	7

Crisis as Catalyst: Why Most Prioritisation Frameworks Are Fair-Weather Tools — and What to Build Instead

Portfolio · No. 69

Originally written August 2009 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“A prioritisation method that cannot stop anything is not a prioritisation method; it is a queue with a scoring column.”

1. Executive Summary

For a decade, organisations invested in ever more elaborate machinery for prioritising their change portfolios: weighted scoring models, strategic alignment matrices, multi-criteria ranking, stage-gate governance. Then the credit crisis arrived, budgets were cut by a third or more in a matter of weeks, and almost none of that machinery proved able to do the one thing it existed for — to stop something. The real triage happened elsewhere: in rooms, under time pressure, on judgement, largely outside the frameworks that were meant to govern exactly this decision.

This paper argues that the crisis did not break portfolio prioritisation; it revealed that most of it had never worked. Conventional scoring frameworks are fair-weather instruments. They are excellent at justifying a portfolio the organisation already wants to fund and useless at forcing the trade-offs it does not want to face, because they are built to rank rather than to constrain. When everything can be funded, ranking is enough. When it cannot, ranking is worthless, because a rank order with no budget line drawn across it stops nothing.

The evidence of the past year points to a specific recommendation: institutionalise the constrained triage the crisis imposed, rather than treating it as an emergency to be survived and forgotten. Concretely, organisations should govern their portfolio permanently against a hard constraint — a forced ranking under a budget deliberately smaller than the one available — so that the discipline of stopping becomes a routine capability rather than a skill exercised only under a gun. The paper weighs this against the alternatives, addresses the real objection that permanent constraint is corrosive, and sets out what building the discipline would require.

2. The prioritisation that only happened under a gun

Through the autumn and winter, a scene repeated itself across organisation after organisation. A portfolio of change — sixty, eighty, a hundred initiatives, funded through the growth years on the comfortable assumption that most good ideas could eventually be afforded — met a budget instruction that arrived without warning and without appeal: reduce spend by a third, and do it this quarter. And the striking thing, in almost every case, was that the organisation's own prioritisation apparatus was not what made the decision.

The weighted scoring model, dusted off for the purpose, returned what it always returned: a list on which nearly everything scored between 3.2 and 3.8 out of 5, clustered so tightly

that it separated almost nothing. The strategic alignment matrix confirmed, reassuringly, that virtually every initiative was aligned to at least one strategic objective – because the objectives were broad enough that almost anything could be mapped to them. Faced with a genuine need to remove a third of the portfolio, the framework offered a ranking in which the thirtieth item was indistinguishable from the thirty-first, which is to say it offered nothing at the only point that mattered.

So the decision was made another way. A small group of senior people sat in a room, argued from knowledge and judgement, and cut. Often they cut well. But they cut despite the governance, not through it, and the fact that the real triage happened outside the framework is the single most important piece of evidence this paper has to offer.

3. Why the frameworks failed when they were finally needed

It would be easy to conclude that the frameworks failed because they were badly built, and to respond by building better ones. That would be the wrong lesson. The frameworks failed because of what they are for, and no amount of refinement changes their purpose. Three structural flaws, present in the method rather than the execution, explain the collapse:

- They rank; they do not constrain. A scoring model produces an ordered list. But an ordered list only forces a decision when a line is drawn across it, and drawing that line – declaring everything below here does not happen – is precisely the act the framework leaves to someone else. In the growth years, no line was ever drawn, so the ranking was never tested. The crisis drew the line, and the ranking had nothing to say about where it should fall.
- They compress rather than separate. Multi-criteria scoring, by averaging many factors, pulls almost everything toward the middle. An initiative that is superb on one dimension and poor on three ends up scoring much like one that is mediocre on all four. The method is, in effect, engineered to make things look similar, which is the opposite of what a triage needs.
- They assume the constraint is soft. Every conventional framework treats the budget as an input to be optimised against, not a wall to be hit. When the wall arrives suddenly and hard, the optimisation logic has no gear for it, because it was designed for a world in which the answer to we cannot afford all of this was always then let us find a little more.

A prioritisation method that cannot stop anything is not a prioritisation method. It is a queue with a scoring column – an ordering of things the organisation intends to do eventually, dressed in the language of choice.

4. The evidence: what the cut actually revealed

The wreckage of the past year is, for once, a rich source of evidence, because the crisis ran the experiment no organisation would ever run deliberately: it forced real stop decisions and let us observe what the frameworks had been hiding. Consider a composite portfolio, representative of what recurred across many.

A change portfolio of roughly seventy initiatives, funded at some tens of millions annually, receives an instruction to remove thirty-five per cent of spend within a single quarter. The scoring model ranks all seventy; the top fifteen and the bottom five are clear enough, but the fifty in the middle score within half a point of one another and cannot be separated by the numbers. When the cuts are finally made by judgement, three findings emerge that the framework had actively obscured:

1. A significant share of the portfolio was already effectively dead. Perhaps a fifth of the initiatives were consuming budget while delivering nothing – stalled, superseded, or continued only because no one had ever been required to stop them. The framework had kept scoring them politely for years. The crisis was the first time anyone asked whether they should exist at all.
2. The highest-scoring initiatives were not the ones protected. When real money was at stake, the room protected initiatives that kept the organisation solvent and compliant – not the ones the strategic-alignment score rated most highly. The revealed priorities and the documented priorities were materially different, which tells you the documented ones had never been load-bearing.
3. Stopping was slower and costlier than expected. Initiatives that should have been simple to cancel turned out to carry contractual commitments, interdependencies, and half-delivered assets that made stopping them nearly as expensive as finishing. The portfolio had no decommissioning discipline because it had never before had to stop anything at scale.

What the framework reported	What the forced cut revealed
All initiatives scored 3.2–3.8, tightly clustered	Roughly a fifth were delivering nothing at all
Highest strategic-alignment scores marked the priorities	The genuinely protected work was solvency and compliance, not the top scorers
Ranking implied a clean cut line	Stopping carried contractual and interdependency costs no score captured
Governance owned the decision	The real triage happened by judgement, outside the framework

The evidence, taken together, is not that the organisations chose badly under pressure. Many chose well. It is that the apparatus contributed almost nothing to the choosing, and that a year of comfortable rankings had concealed a portfolio in which a fifth of the spend was waste hiding in plain sight.

5. What genuinely worked under constraint

If the frameworks failed, something else succeeded, and it is worth being precise about what, because the recommendation depends on it. What worked, in the rooms where triage went well, had three features that the scoring models lacked. First, a hard constraint stated up front – not rank these but this is the money; what survives? The constraint did the separating work the scores could not, because a fixed budget forces genuine comparison in a way an open-ended ranking never does. Second, forced trade-off – the discipline that funding one thing meant naming the thing it displaced, which turned prioritisation from an additive exercise into a genuinely subtractive one. Third, judgement held accountable – named people owning the cut and defending it, rather than a score

absorbing the responsibility no individual wished to carry.

None of this required a new framework. It required a different question. The scoring models asked how good is this initiative? The successful triage asked what would we give up to keep it? – and that second question, asked under a real constraint, is the whole of prioritisation. Everything else is decoration.

6. The options: how to keep what the crisis taught

The temptation, as the immediate pressure eases, will be to file the past year under emergency and return to the comfortable machinery. That would waste the most expensive lesson the organisation has been handed in a decade. There are, realistically, three options for what to do instead, and honesty requires weighing them rather than leaping to the preferred one.

- Option one: return to the frameworks, refined. Keep the scoring models but improve them – better criteria, sharper weighting, more discipline in application. This is the least disruptive path and the easiest to sell, because it changes nothing fundamental. Its fatal weakness is that it addresses execution when the failure was structural: a better-tuned ranking still stops nothing, because ranking is not the missing capability.
- Option two: rely on judgement, invoked when needed. Accept that the frameworks are theatre and simply convene the room when a real cut is required, as the crisis did. This is honest about where decisions actually get made, but it makes the discipline hostage to crisis – available only when fear is high enough to convene it, absent in the calm years when portfolios quietly refill with waste. It treats the muscle as something to be borrowed in emergencies rather than kept in condition.
- Option three: institutionalise the constraint. Govern the portfolio permanently against a hard budget deliberately set below the funds available – running the standing question if we had a third less, what would survive? as routine practice rather than crisis response. This keeps the subtractive discipline alive in calm conditions, surfaces the dead initiatives continuously rather than once a decade, and builds stopping into a capability the organisation actually possesses.

The third option is harder than the first and more durable than the second, and it carries a real cost that its advocates should not hide.

7. The objection worth taking seriously

The strongest case against institutionalising the constraint is that permanent crisis discipline is corrosive. Run an organisation perpetually as though the money were about to run out, the argument goes, and you exhaust its people, starve the long-horizon bets that never survive a constrained ranking, and eventually cry wolf so often that the constraint loses the very force that made it work in the crisis. There is truth in this, and it would be dishonest to wave it away: a constraint that is always maximally tight does become background noise, and an organisation that can only ever subtract will, in time, forget how to invest.

The answer is not to dismiss the objection but to design against it. The institutionalised constraint should be firm, not brutal – a standing discipline of forced trade-off, not a permanent austerity. Its purpose is to keep the portfolio honest and the muscle of

stopping in use, not to grind spending down year on year. And it should explicitly protect a small, ring-fenced allocation for long-horizon work that the annual ranking is not permitted to touch, precisely so that the discipline of subtraction does not quietly amputate the future. Constrained is not the same as starved, and the design has to make the difference real.

8. Recommendation

The recommendation of this paper is the third option, designed against its own worst tendency. Organisations should institutionalise constrained prioritisation — governing the change portfolio against a hard, deliberately sub-maximal budget; requiring every funded initiative to name what it displaces; making judgement accountable rather than hiding it behind a score; and ring-fencing a protected allocation for the long-horizon work that constraint would otherwise kill. The scoring frameworks need not be abolished, but they should be demoted to what they are good at — informing a judgement — and stripped of the authority they never earned, which was to substitute for one.

“The organisations that emerge strongest will not be the ones that survived the cut. They will be the ones that never gave back the discipline the cut forced on them.”

9. What it would take

Building this is less a matter of new tooling than of new habit, and the requirements are modest but non-negotiable. It takes a governance forum with the standing authority to stop things, not merely to rank them. It takes a portfolio held continuously against a stated constraint rather than reset to comfort at each annual round. It takes the decommissioning discipline the crisis found missing — the ability to stop an initiative cleanly, contracts and interdependencies and half-built assets included, rather than discovering at the moment of cancellation that stopping was never designed for. And it takes leaders willing to keep asking the uncomfortable subtractive question in the calm years, when nothing external compels it and every incentive pulls back toward the easy additive ranking.

That last requirement is the binding one. The frameworks did not fail for want of sophistication; they failed because no one ever had to draw the line until the crisis drew it for them. The whole of the recommendation reduces to keeping the hand that draws the line in practice — so that the next time the constraint arrives, whether from a downturn or from ordinary good sense, the organisation meets it with a capability rather than a room full of people improvising in the dark.