

BEYOND SCORING THEATRE

Strategic Alignment Scoring — The Art of Retrospective Justification

How portfolio prioritisation became a ritual of confirmation rather than a discipline of choice



Strategic Alignment Scoring – The Art of Retrospective Justification

Portfolio · No. 51

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

“The scoring framework is the tribute that politics pays to rationality.”

The Scoring Ritual

The quarterly portfolio review has a familiar choreography. The scoring templates arrive – usually a spreadsheet, occasionally something more sophisticated – and project sponsors consult the strategic objectives approved at the last strategy offsite. Five or six broad themes, printed on a laminated card that sits in a desk drawer between reviews. Each sponsor rates their initiative against each objective. The numbers are assembled into a composite score, ranked, and presented to the investment committee in a neatly formatted slide deck. The committee approves largely what was already in motion. The organisation moves on, satisfied that it has exercised strategic discipline.

We have all participated in this ritual. And most of us know, if we are honest with ourselves, that it is a ritual – not a decision.

The pattern I have observed across financial services, telecommunications, and public-sector organisations is remarkably consistent. The strategic alignment score is not used to determine which initiatives proceed. It is constructed after the decision has already been made – through corridor conversations, executive sponsorship, and the gravitational pull of projects that are already consuming budget. The score exists to provide documentary evidence that the decision was rational.

This is scoring theatre. And the damage it does is not trivial.

The Mechanism of Retrospective Justification

The process works because the scoring criteria are almost always broad enough to accommodate any reasonably articulated initiative. When an organisation's strategic objectives include phrases such as "enhance operational efficiency," "improve customer experience," or "build organisational capability," it would take a peculiarly narrow project to fail to claim alignment with at least two of them. A data warehouse migration scores well on efficiency. So does an office relocation. So does a staff restructuring programme. The strategic objectives are written to inspire, not to discriminate – and scoring against them inherits that same permissiveness.

The real decision – whether to fund this initiative rather than that one – happens elsewhere. It happens when the Group Finance Director mentions over coffee that the payments platform is becoming a regulatory risk. It happens when the new Head of Operations arrives from a competitor and brings a transformation agenda built around her previous organisation's model. It happens when a programme that has consumed eighteen months and forty million pounds cannot be stopped because the sunk cost has become

politically unmentionable.

The scoring exercise runs parallel to these forces, not through them. It arrives after the constellation of power has already shaped the portfolio, and its function is to provide a veneer of analytical rigour over decisions that are fundamentally political.

Consider a pattern that will be familiar to anyone who has sat on a portfolio board. A project sponsor submits an initiative with a composite alignment score of 3.8 out of 5. A competing initiative scores 3.6. The investment committee approves the 3.8 – not because the difference of 0.2 is analytically meaningful, but because the sponsor of the 3.8 is a member of the executive team and the sponsor of the 3.6 is a middle manager in a regional office. The scoring has not made the decision. It has dressed the decision in the language of methodology.

Why the Pattern Persists

If the scoring is theatre, why does every organisation persist with it? Because it serves two purposes that have nothing to do with strategic prioritisation.

The first is governance compliance. Most mature organisations – and certainly those in regulated sectors operating under the weight of Sarbanes-Oxley or FSA expectations – need documentary evidence that investment decisions followed a structured process. The alignment scoring provides that evidence. It creates an auditable record showing the strategic rationale for each approved initiative, even when the rationale was constructed after the fact. For the governance function, the score's accuracy matters less than its existence.

The second is political protection. In any organisation of reasonable size, the portfolio review is a contest for resources. The strategic alignment score gives every sponsor a defensible position: "My project scored 4.2 out of 5 on strategic alignment." It does not matter that the competitor project scored 4.1, or that the difference between those scores is analytically meaningless. The number provides cover. It transforms an argument about power and priority into an argument about methodology – which is a much safer argument to have in a committee room.

These two functions – compliance and protection – are powerful enough to sustain the practice indefinitely, regardless of whether it produces better portfolio decisions. The scoring framework becomes self-perpetuating: too embedded in the governance architecture to remove, and too politically useful to challenge.

What Genuine Alignment Would Require

The problem is not that organisations attempt to align their portfolios with strategy. The problem is that they confuse a scoring exercise with a genuine alignment discipline – and the difference between the two is uncomfortable.

Genuine strategic alignment would require, at minimum, three things that most organisations are unwilling to do.

First, it would require strategy that discriminates. An organisation whose strategic objectives can accommodate any initiative has not made strategic choices – it has produced a list of aspirations. Real alignment scoring demands strategic objectives

specific enough that a significant proportion of proposed initiatives genuinely fail to meet them. This is politically difficult because broad objectives create the illusion of consensus, and narrowing them forces senior leaders to acknowledge that some of their priorities are not, in fact, strategic.

Second, it would require scoring before commitment. The temporal sequence matters enormously. If the alignment assessment happens after the project has been approved, funded, or – worse – started, then it is retrospective justification by definition. Genuine alignment means the scoring shapes the decision, which means it must happen before political commitments have been made. In practice, this would mean rejecting or deferring initiatives that have already been promised to business units – something few portfolio governance structures are designed, or authorised, to do.

Third, it would require willingness to defund. The ultimate test of a prioritisation mechanism is whether it can stop things. An alignment scoring process that has never led to the cancellation of a funded initiative is not a prioritisation process; it is a documentation process. The distinction matters because it determines whether the organisation's resources are genuinely directed by strategy or merely annotated with strategic labels after the fact.

The test of any portfolio prioritisation mechanism is not whether it can rank things. It is whether it has ever stopped one.

One might object that demanding this level of rigour is unrealistic – that organisations are political systems, that perfect rationality in resource allocation is a fiction, and that the scoring framework at least imposes some structure on what would otherwise be a free-for-all. This objection has merit, as far as it goes. Organisational decisions are never purely rational, and some structure is better than none. But the objection mistakes the nature of the criticism. The problem is not that scoring fails to achieve perfect rationality. The problem is that it claims to achieve what it does not – and in doing so, it removes the pressure to build something that actually works.

The Honest Alternative

We are left, then, with an uncomfortable professional reality. The strategic alignment scoring frameworks we build, maintain, and present are, in the majority of cases, instruments of confirmation rather than instruments of choice. We know this – most experienced portfolio practitioners will acknowledge it privately – and yet we continue to operate them because the governance architecture demands it and because challenging the practice means challenging the political economy of investment decisions.

The honest response is not to abandon alignment scoring. The governance need is real, and the alternative – undocumented, purely political allocation – is worse. The honest response is to stop pretending that the current practice does what it claims to do.

This means designing scoring frameworks that are genuinely discriminating, knowing that most organisations will resist. It means insisting that the scoring precedes the commitment, knowing that it usually follows. It means measuring the framework's value not by the elegance of its methodology but by whether it has ever changed a decision. And it means being candid with our boards and steering committees about the difference between a scoring exercise and a strategic discipline – rather than allowing the former to masquerade as the latter.

The scoring framework is the tribute that politics pays to rationality. We can accept that and work to narrow the gap, or we can continue to mistake the tribute for the thing itself. The profession has been doing the latter for long enough.