

BEYOND THE BENEFITS CASE

The Portfolio That Cannot Stop: An Honest Account of Outcome-Based Measurement

A Practitioner's Perspective on Why Portfolio Outcome Measurement Works Only When It
Is Allowed to Hurt



The Portfolio That Cannot Stop: An Honest Account of Outcome-Based Measurement

Portfolio · No. 41

Originally written December 2007 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“Timing, not precision, is what makes an outcome measure actionable.”

The Version Everyone Approves Of

There is a version of outcome-based measurement that everyone approves of and nobody fears. It lives in the benefits column of a business case, is reported each quarter in the portfolio pack, and confirms, cycle after cycle, that the investments the organisation has already decided to make are delivering the value it already assumed they would. It is sincere, well-intentioned, and almost entirely useless. I have watched portfolios run this way for years, and the honest account is that measuring outcomes changed nothing — because nothing was ever going to change as a result of what the measures said.

Outcome-based measurement that actually works looks different, and it is uncomfortable in a way the approved version never is. In a sentence, what it looks like is a portfolio that stops things. Not one that reports on things, or re-forecasts things, or amber-rates things — one that takes an initiative it has already spent real money on, looks honestly at whether the outcome is arriving, concludes that it is not and will not, and cancels it.

Until a measurement system has stopped an initiative in public, it has not been tested — and the organisation does not yet know whether it has one.

The One Job

The reason this is the true test is that outcome measurement, at portfolio level, has only one job the older machinery could not already do. Cost and schedule were always measured. Activity was always measured. What benefits realisation was invented to add was the ability to look at money already committed and say: this is not working, and continuing to fund it is a worse decision than stopping. Everything else outcome measurement produces — the dashboards, the benefit maps, the value scores — is instrumental to that one decision. Strip the decision away and the rest is decoration, however elaborate.

So the honest account begins with an admission: most portfolios that claim to do outcome-based measurement have never once stopped an initiative because of it. They stop initiatives, certainly, but for other reasons — a budget cut, a change of sponsor, a reorganisation. The outcome measures are consulted after the fact to justify a decision taken on other grounds, not before it to make one. This is measurement as ceremony, and it is the natural resting state of any portfolio that has not deliberately chosen otherwise, because every force in the organisation pushes toward it.

Why Stopping Is So Hard

Consider what those forces are. An initiative that has been running for two years has a sponsor whose credibility is bound up in it, a team whose jobs depend on it, and a business case whose benefits have already been banked in someone's plan. To cancel it on the evidence of an outcome measure is to override all of that with a number, and numbers are weak against reputations. The sunk cost is felt as a reason to continue rather than, as the textbook insists, an irrelevance. And the person who proposes the cancellation is proposing to be the one who was wrong — or who made someone else wrong — in a room full of people who would rather the question had not been asked.

Against those forces, a scorecard is nothing. This is why the organisations where outcome measurement genuinely works are not distinguished by better metrics; their metrics are often cruder than the ones on display in portfolios where nothing is ever stopped. What distinguishes them is that they have made stopping survivable. A few things, in my experience, make the difference.

1. They separate the decision to stop from the judgement of the people involved. Cancelling an initiative because its outcome is not arriving is treated as new information acted upon, not as a failure to be punished. Where stopping ends careers, stopping stops happening, and the measures learn to always say continue.
2. They measure outcomes early enough to stop cheaply. A portfolio that only sees outcome data at the end of an initiative can conduct only post-mortems. The ones that act watch the leading signals — that adoption is not coming, that the behaviour change the benefits depend on is not materialising — while the amount still to be spent is larger than the amount already sunk.
3. Someone senior has to kill their own initiative on the evidence — visibly, and more than once. Until that has happened, everyone correctly reads the measurement system as a thing that applies to other people's projects, and the ceremony continues. The tone of a portfolio is set not by its framework but by what its most powerful people do when the measures turn against something they own.

The Narrow, Harder Point

I am not arguing that portfolios should become trigger-happy, cancelling on the first disappointing data point. Outcomes are noisy and often lag the work that produces them; a measure that falls short may mean the initiative is failing, or merely that it is not yet finished, and telling those two apart is most of the skill. The point is narrower and harder. A portfolio that has structurally deprived itself of the ability to stop — through sunk-cost thinking, through punishing the messenger, through measuring too late to act — has not implemented outcome-based measurement at all, whatever its dashboards suggest. It has implemented a more sophisticated way of confirming decisions it was always going to make.

What outcome-based measurement looks like when it works, then, is not a cleaner report. It is a slightly uncomfortable portfolio review, held often enough to matter, in which the genuinely live question is whether to keep funding each thing — and in which the answer, often enough to be believed, is no. Everything else is the approved version: sincere, tidy, safe, and changing nothing at all.