PORTFOLIO BOOK SERIES · No. 94

Perspective

BEYOND THE TWO-BY-TWO

Why the Matrix Never Tells You What to Stop

Portfolio prioritisation is treated as a ranking problem. On the delivery side it is a stopping problem — and no quadrant has ever closed an initiative.

Giovanni Leonardi

August 2024

Why the Matrix Never Tells You What to Stop

Portfolio · No. 94
Originally written August 2024 | Re-edited and re-rendered for the WhereBeyond Edition,

> *"A matrix can tell you, quite well, which things are worth doing. It has never once told an organisation what to stop."*

1. The review that ends where it began

Forty-seven initiatives on a single slide. Each one a neat dot on a two-by-two grid — value climbing the vertical axis, effort running along the horizontal — and almost every dot clustered in the same corner: high value, low effort, the quadrant everyone wants to live in. The room has spent six weeks preparing. Business cases were written, scores were debated line by line, a facilitator was brought in to keep things moving. And after ninety minutes the portfolio that leaves the room is the portfolio that entered it: forty-seven initiatives, now wearing a ranking.

Everyone leaves satisfied that the process was rigorous. Nobody has agreed to stop a single thing.

The pattern recurs across almost every prioritisation exercise I have watched from the delivery side, and its shape is fixed. The matrix is articulate about what deserves to be done and mute on the one question that matters once an organisation is already full: which of the things currently in flight should now cease. A ranking answers what is most worth starting. It does not answer what must we put down to make room — and on the delivery side, that second question is the entire game.

2. A ranking instrument asked to perform a subtraction

The two-by-two is, at heart, a sorting device. Give it a set of candidate items and a pair of axes and it will arrange them from most to least attractive. That is a genuinely useful thing to do when you are choosing among options you have not yet committed to. The trouble is that portfolio prioritisation in a real organisation is almost never a green-field choice among fresh candidates. It is a negotiation with a portfolio that is already running, already staffed, already promised to sponsors who spoke at the last town hall.

So the matrix gets asked to do something it was never built for. Ranking is an additive operation: it tells you the order in which to pick things up. Prioritisation, as the delivery side lives it, is a subtractive one: it is the decision about what to put down. These are not two views of the same act. Picking up and putting down have different mechanics, different politics, and — crucially — different quadrants, except that the matrix only ever draws the ones for picking up.

> *A ranking tells you the order in which to start. It says nothing about the order in which to stop — and an organisation that is already full can only start by stopping.*

Watch what actually happens to the initiatives that land in the bottom-right of the grid, the low-value, high-effort quadrant that theory says should die. In a green-field choice they would simply never be selected. But these are not green-field. Each of those unattractive dots already has a team, a board of open tickets, a partly-built integration, a director who defended its funding in March. The matrix has ruled against it. And the following Monday it is still being worked on, because a dot moving into a bad quadrant is not the same event as a person being told to stop.

---

3. Why the quadrants cannot see what is already running

The deeper reason the matrix never tells you what to stop is that its axes are blind to the two facts that govern stopping: sunk commitment and shared capacity.

Consider a portfolio I will keep deliberately composite. Forty-seven funded initiatives. A single integration platform team — nine engineers — sitting underneath thirty-one of them as a hard dependency. On the matrix, all forty-seven are scored independently, each on its own value and its own effort, as though the nine engineers could be in thirty-one places at once. The grid has no axis for contention. It cannot show that initiative number twelve, sitting proudly in the top-left, is only "low effort" on paper because its effort is quietly borrowed from a team that thirty other initiatives are also counting on.

The number that exposes this is never on the slide. The nine-person team can, realistically, carry perhaps four of those thirty-one initiatives at a time without the queue collapsing into thrash. So the true prioritisation decision is not "rank forty-seven by attractiveness." It is "choose the four, and consciously stop or park the other twenty-seven that are currently pulling on the same nine people." The matrix cannot even represent that sentence. It has no cell for the initiative that is individually worthwhile and collectively unaffordable, which is the exact description of most of what a full portfolio is drowning in.

- The matrix scores initiatives in isolation; delivery is constrained by what they share.
- It scores expected value; it ignores the sunk commitment that makes stopping feel like a loss.
- It scores effort as a fixed quantity; it cannot see effort as a claim on a contended team.
- It produces an order; the organisation needs a line, and then the will to enforce it.

Anchor this in where the discipline actually stands. Much of the better thinking has already moved this way — the shift from funding projects to funding persistent product teams, capacity-based planning, the language of cost of delay and weighted-shortest-job-first, the whole "project to product" argument for stable teams that pull work rather than having work pushed onto borrowed people. All of it is, underneath, an admission of the same truth: that the binding constraint is the standing capacity of teams, not the attractiveness of ideas. And yet the annual prioritisation ritual still opens with the two-by-two, because the two-by-two is easy to build and comfortable to present, and because it lets a leadership team feel decisive without anyone having to say the word "stop" out loud to a colleague whose initiative is being stopped.

4. "But a forced ranking does force stops"

The strongest objection to all this deserves a fair hearing, because on its own terms it is not wrong. A disciplined portfolio team, the argument goes, does not stop at the picture. It takes the ranked list, draws a funding line across it based on real capacity, and everything below the line is by definition cut. Force-rank to a single ordinal list, cap it at what the organisation can actually staff, and the matrix has told you what to stop: everything beneath the line. The tool is fine; weak organisations simply fail to enforce it.

I have watched serious teams do exactly this, and it is a real improvement over the decorated cluster. But it mistakes drawing the line for enforcing it. The line gets drawn on the slide. What does not happen — almost ever — is the second act: walking to the twenty-seven initiatives below the line and de-staffing them. The people on a below-the-line initiative do not evaporate when the line is drawn. Their director was in the room and negotiated an exception "just to finish the current phase." The initiative is seventy percent done and stopping it feels like conceding the first seventy percent was waste. So the initiative keeps breathing — throttled, under-resourced, quietly stealing a day a week from the very engineers the top-four initiatives are depending on. The line existed on paper. It was never cut into the organisation.

> *"Everything below the line kept breathing — throttled, under-resourced, and quietly stealing a day a week from the teams above it."*

That is the failure mode the matrix cannot fix, because it is not a scoring failure. The scores were fine. It is that the instrument ends its work at the moment the hard work begins. Ranking is complete when the list is ordered. Stopping is complete only when a named team is doing something different on Monday — and nothing in a two-by-two, or in the forced ranking built on top of it, reaches that far.

5. Prioritisation is de-staffing, or it is theatre

If there is one conviction the delivery side earns, paid for in stalled roadmaps and half-built platforms, it is this: prioritisation that does not change who is working on what has not happened. It has merely been simulated. A prioritisation exercise should be measured by exactly one output — the list of things that were being done last week and are not being done this week, and the names of the people who have been moved off them. If that list is empty, the exercise was a costume drama, however elegant the grid.

This reframes what a portfolio review is for. It is not an occasion to admire a field of dots. It is the one moment an organisation has to perform subtraction deliberately rather than letting it happen by accident — because subtraction happens either way. In a full portfolio, starting the top-four always stops something; the only question is whether the something is chosen in the room or defaulted to in the corridor, where the initiative that gets starved is not the least valuable but the least politically defended.

6. What the delivery side would put in its place

None of this is an argument against prioritising. It is an argument for prioritising the constraint instead of the candidates. A few disciplines carry more weight, from where delivery sits, than any refinement of the axes:

1. Start from capacity, not from candidates. Name the genuinely contended teams first — the nine-person platform group, the two people who understand the pricing engine — and treat their throughput as the fixed budget. Everything else is a claim against that budget.
2. Make the stop-list the primary output. Do not leave the room with a ranked list of what to do. Leave with the explicit, named list of what stops, and who moves off it,

dated for this week. If nothing stops, nothing was prioritised.

3. Price the sunk cost honestly and then ignore it. Say plainly that an initiative is seventy percent done, acknowledge the discomfort, and decide on the value of the remaining thirty percent against everything else that team could do instead. The first seventy percent is spent whatever you choose.

4. Govern the line after the meeting, not in it. The exceptions that dissolve a prioritisation decision are all requested after the slide is approved. Someone has to own the unglamorous work of refusing them. That ownership, not the scoring model, is where prioritisation actually lives or dies.

The two-by-two will survive all of this, and it should; as a way of opening a conversation about relative worth it has its place. The error is asking a picture of attractiveness to carry a decision about subtraction. A matrix can tell you, quite well, which things are worth doing. It has never once told an organisation what to stop — and until a portfolio review is judged by what it stopped rather than by how tidily it ranked, the same forty-seven initiatives will keep walking out of the room they walked into.