

BEYOND THE SURVIVOR LIST

The Investment Loop

One Method, Any Scale

Giovanni Leonardi

2026



Contents

Executive Summary 6

The problem this book solves	7
The central model	7
The two jobs	8
The unit and the promise	8
One method, three settings	8
What is in this book	8
How to read this book	9

Part I — Foundations 10

1. Why Portfolios Drift 11

1.1 The portfolio nobody chose	12
1.1.1 How the residue forms	12
1.2 What drift costs	13
1.3 The tells	14
1.4 Doing the right things	15
1.5 The promise	15

2. The Canon 17

2.1 A pool, not a list	18
2.2 The unit: the investment	18
2.3 The two jobs	19
2.4 The Investment Loop	20
2.4.1 Which stage does which job	21
2.5 The two-way link with strategy	21
2.6 Two kinds of measures	22
2.7 One method, three settings	22
2.8 What the rest of the book is	24

Part II — Working the Investment Loop 25

3. Direction 26

3.1 Where strategy becomes machinery	27
3.2 The reference framework	27
3.3 Strategy sets the weights, not the scorecard	28
3.4 Risk appetite	29
3.5 Funding buckets	30
3.6 Direction is where the loop closes	31

3.7 Direction at three settings	32
4. Demand	34
4.1 The visibility problem	35
4.2 What makes a candidate an investment	35
4.3 In-flight work enters the same door	36
4.4 The intake log and the single front door	37
4.5 Flagging uncertain bets	38
4.6 Demand at three settings	38
5. How We Value	40
5.1 The comparison problem	41
5.2 The models and when to use them	41
5.2.1 Weighted scoring against the reference framework	41
5.2.2 Cost of Delay and WSJF	42
5.2.3 Value-versus-effort (the simple portfolio map)	42
5.2.4 Financial appraisal — ROI, NPV, payback	43
5.3 Which model for which setting	43
5.4 Applying the strategy-set weights	44
5.5 Forward value: the only number that matters for the decision	44
5.6 Valuing uncertain bets as options	45
5.7 How performance data updates the value	46
5.8 How We Value at three settings	46
6. How We Choose the Mix	48
6.1 From values to a portfolio	49
6.2 The funding line and what it means	49
6.3 Capacity versus demand	50
6.4 Balance across four dimensions	50
6.4.1 Risk and return	51
6.4.2 Time horizon	51
6.4.3 Strategic coverage	51
6.4.4 Run-the-business versus change-the-business	52
6.5 Concentration risk	52
6.6 Funding buckets in action	53
6.7 Making real trade-offs	53
6.8 How We Choose the Mix at three settings	54
7. Fund	55
7.1 From decision to commitment	56
7.2 Modes of funding	56
7.3 Cadence: the heartbeat of the portfolio	57
7.4 Staged funding for uncertain bets: fund-to-learn	58

7.5 The funding and commitment sheet	59
7.6 Fund at three settings	59
8. Steer	61
8.1 The second job begins	62
8.2 What Steer is for	62
8.3 The anatomy of a Steer problem	63
8.3.1 Dependencies	63
8.3.2 Scarce resource contention	63
8.3.3 Sequencing	64
8.3.4 Duplication	64
8.3.5 Blockers the investment cannot clear alone	64
8.4 Steer forces decisions	65
8.5 Feeding hard truths to Review	65
8.6 The Steer toolkit	66
8.7 Steer at three settings	67
9. Continuous Re-prioritisation and Reallocation	68
9.1 The loop comes back around	69
9.2 Re-judging the live set	69
9.3 Moving money on forward value	70
9.4 Tracking benefits after delivery	71
9.5 Stopping work	72
9.6 The politics and the biases	72
9.7 Designing for survivable decisions	73
9.8 Benefits after delivery and the feedback to strategy	74
9.9 Continuous re-prioritisation at three settings	74
Part III — Running It For Real	76
10. The Operating Model	77
10.1 The method needs a home	78
10.2 Roles	78
10.3 Decision rights	79
10.4 Governing the prioritisation process itself	80
10.5 Setting the dials	80
10.6 Portfolios of portfolios	81
10.7 How the dials evolve	82
11. Making It Stick	83
11.1 The method meets the organisation	84
11.2 Communicating the decisions	84
11.3 Carrying people with you	85

11.4 The usual failure modes	85
11.5 Data traps	87
11.6 Making the rules credible	87
11.7 Sustaining the rhythm	88

Annex A — Direction — Templates 89

About these templates	89
A1 — Strategy-on-a-page	89
A2 — Investment criteria and weighting framework (the reference framework)	90
A3 — Risk-appetite sheet	91
A4 — Allocation-buckets framework	91

Annex B — Demand — Templates 93

About these templates	93
B1 — The Investment One-Pager	93
Section 1: What and why	93
Section 2: Strategic fit	94
Section 3: Cost and resource	94
Section 4: Risk and confidence	94
Section 5: Forward value (in-flight only)	94
B2 — Intake and demand log	95

Annex C — Value & Decide — Templates 96

About these templates	96
C1 — Scoring model	96
C2 — Prioritisation-model selector	97
C3 — Portfolio balance and risk-return view	97
C4 — Capacity-versus-demand view	98

Annex D — Fund — Templates 100

About these templates	100
D1 — Funding and commitment sheet	100
D2 — Funding cadence calendar	101
D3 — Staged fund-to-learn card	101
Tranche structure	102

Annex E — Steer — Templates 103

About these templates	103
E1 — Portfolio dashboard and health view	103
E2 — Dependency map	104
E3 — Decision and escalation log	104

E4 — Resource-contention view	105
---	-----

Annex F — Review — Templates 107

About these templates	107
F1 — Value and benefits tracker	107
Part 1: In-flight value tracking	107
Part 2: Post-delivery benefit tracking	107
F2 — Stop-or-continue template	108
F3 — Lessons log	109

Annex G — Cross-cutting — Templates 111

About these templates	111
G1 — Roles and decision-rights matrix	111
G2 — The Dial-Setting Worksheet	112
Step 1: Describe the current reality	112
Step 2: Score the dials	112
Step 3: Select the setting	113
G3 — Stakeholder and communication map	113
G4 — Portfolio health and risk-balance card	114

Executive Summary

Executive Summary

The problem this book solves

Most portfolios are not chosen. They are accumulated — the residue of everything ever started and never stopped, defended by inertia and sustained by the absence of a systematic way to say no. The result is familiar: resources spread too thin across too many things, zombie initiatives consuming budget without producing value, a strategy that says one thing and a funded set that does another, and a permanent shortage of room for the new investments that could change what is possible. The deeper problem — the one that makes it worse — is that the best work that was never funded remains invisible. Nobody knows what the organisation could have done with the money it spent on the wrong things.

This book is about fixing that. Not through a one-off rationalisation — organisations have done those, and the portfolio drifts straight back — but through a method that makes continuous, deliberate re-choosing structural. A portfolio managed well is not one that is decided once and executed faithfully. It is one that is re-decided, on a rhythm, with the money always moving to follow the value wherever it now is.

The central model

The method is built on The Investment Loop: six continuous stages that feed back into one another.

Direction turns strategy into the machinery of choice — the selection criteria, their strategy-set weights, the risk appetite, and the funding buckets that pre-commit the pool across categories. Demand surfaces every candidate, new and in-flight, into a single comparable pool. Nothing competes off the books. Decide values each candidate using the right prioritisation model for its type and setting, and chooses the mix the pool can actually afford — balancing across risk, time horizon, strategic coverage, and concentration. Fund releases the money on a rhythm, with uncertain investments funded in staged tranches to learn rather than deliver. Steer makes the chosen mix work together while it runs — managing dependencies, resolving resource contention, sequencing investments, and clearing blockers that no single owner can clear alone. Review re-opens the whole question: re-judges the live set against new demand on the same criteria, moves money to where the forward value now is, tracks whether benefits actually arrived after delivery, and stops what no longer earns its place.

The loop is not a line. Review feeds back into Direction and Demand. The portfolio receives strategy and sends evidence back: what worked, what assumptions proved wrong, what the criteria have been selecting. Both directions of that link are part of the method.

The two jobs

A portfolio has two jobs, and they are genuinely different. The first is to choose the right mix of investments and to keep re-choosing it — the subject of Direction, Demand, Decide, Fund, and Review. The second is to make the chosen mix work together — the subject of Steer. A perfectly chosen mix, left to collide with itself, will still fail. The portfolio's responsibility ends at the boundary between investments, not inside any one of them: coordination between investments is the portfolio's work; delivery within an investment belongs to the investment.

The unit and the promise

The unit of management is the investment — one thing the portfolio chooses to fund for a return. Programme, project, product and initiative are all types of investment. Using one unit creates one queue and one set of questions, which is what makes the method work at every scale.

The core promise is this: money already spent gets no vote. Every funding decision is made on forward value — value still to come against cost still to spend. A running investment is not protected by the fact of having started; it earns continuation by offering more forward value than the alternatives. This is the principle that breaks the cycle of drift — and it is the one that requires the most discipline to keep.

One method, three settings

The method scales through a single mechanism: one set of questions, adjustable by four dials — cadence, who decides, formality, funding rigidity. The dials resolve into three named settings: Lean for a small, fast portfolio run by one or a few people; Managed for a portfolio that needs a standing board and a structured cycle; Enterprise for a large, formally governed portfolio operating through tiered boards. The questions never change. Only the speed, the number of people in the room, and the weight of the paperwork differ. An organisation can run different parts of itself at different settings, and can change its own setting as it grows, without abandoning the method.

What is in this book

Part I — Foundations — establishes the problem and the vocabulary. Chapter 1 diagnoses the drifted portfolio and its costs. Chapter 2 sets the canon: the Investment Loop, the two jobs, the three settings, and the two kinds of measures (selection criteria, which choose; and performance measures, which track — they are not the same, and conflating them is one of the most expensive confusions in portfolio management).

Part II — Working the Investment Loop — takes each stage in turn, with Lean, Managed, and Enterprise variations shown inside each stage rather than in separate sections. The

valuation chapter covers weighted scoring, Cost of Delay, value-versus-effort, and financial appraisal — with a model selector and a clear account of how to value uncertain bets as options rather than fixed returns. The choosing chapter covers thresholds, capacity, balance, and concentration risk. The Steer chapter and the Review chapter are the two heaviest, because they carry the most demanding work: the operational coordination that makes the mix function, and the periodic re-decision that keeps the money following the value.

Part III — Running It For Real — covers the operating model (roles, decision rights, the dials, portfolios-of-portfolios, and how the dials evolve with maturity) and the discipline of making the method stick against the forces — political, psychological, and structural — that will otherwise erode it.

Seven annexes map templates to stages: Direction, Demand, Value and Decide, Fund, Steer, Review, and cross-cutting. The templates are skeletons, scaled to the Managed setting, with adaptation notes for Lean and Enterprise. A companion volume — Book 2, the toolkit — will provide fuller self-contained tool cards for each template.

How to read this book

The reader for whom this book is written is the person accountable for a portfolio's choices: a portfolio director, a PMO or EPMO lead, or an agile portfolio owner. The secondary readers are the teams who run the mechanics and the executives who sit on governance.

The book is designed to be read in order: the canon chapter establishes the language that every subsequent chapter uses. Readers who are already familiar with portfolio theory may begin at Chapter 3, but should read Chapter 2 first to confirm that the terminology aligns with what follows.

The method described here is Book 1 — the methodology. It is complete and self-standing. The tools are described functionally in the chapters and provided as working templates in the annexes. The companion Book 2 will deliver each tool as a fully self-contained card, with worked examples and fuller guidance for practitioners building the instruments for the first time. Readers who want to start using the method now have everything they need here to do so.

PART

Part I — Foundations

CHAPTER 1

Why Portfolios Drift

Why Portfolios Drift

1.1 The portfolio nobody chose

Walk into almost any organisation and ask to see its portfolio. You will be handed a list. It will be long. It will be sorted by something — start date, sponsor, business unit, the colour of the status light — and somewhere on it will be every initiative that anyone, at any point in the last several years, managed to get funded and has not yet been forced to admit is finished.

That list is not a portfolio. It is a sediment.

Nobody sat down and chose it. No single decision produced it. It is the accumulated residue of hundreds of small, locally reasonable approvals, each made when money was a little looser, each defended ever since by the simple fact of already existing. The organisation did not decide to run these forty things rather than some better forty. It decided to run the first one, and then the next, and then the next, and the rest is just what survived.

This is the central problem of portfolio life, and it is worth being precise about, because it does not look like a problem. It looks like activity. The teams are busy. The reports are produced. The governance forum meets on the second Tuesday of the month and approves the minutes of the previous second Tuesday. Everything is moving. The trouble is that nobody can tell you, with a straight face, that this particular set of moving things is the best set the organisation could be funding with the money it has. They can only tell you that it is the set that happens to exist.

“A portfolio that nobody chose is not neutral. It is a series of decisions made by inertia and defended by sunk cost.”

1.1.1 How the residue forms

Drift is not a failure of intelligence. It is the natural result of two habits that, on their own, look like virtues.

The first is additive intake. New ideas arrive constantly — from the strategy refresh, from the regulator, from the loudest customer, from the executive who came back from a conference with a thought. Each is assessed on its own merits, in its own meeting, against the question is this worth doing? And very often the honest answer is yes, it is. So it is added. The portfolio grows by accretion, one defensible yes at a time. What is almost never asked, at the moment of the yes, is the only question that actually matters at

portfolio level: is this worth doing more than the thing it will take resources from? That question has no natural owner in an additive process, so it does not get asked, and the list gets longer.

The second habit is the absence of an exit. Starting work is an event. It has a sponsor, a business case, a launch, a kickoff with pastries. Stopping work is a non-event — or worse, an admission. There is no ceremony for it, no template, no person whose job it is to propose it. Initiatives therefore do not leave the portfolio when they stop earning their place; they leave it only when they are completely finished, or when someone with enough authority and enough courage decides to take the reputational hit of killing something. Since finishing is rare and courage is rationed, almost nothing ever leaves. The portfolio is a venue you can check into but, in practice, cannot check out of.

Put additive intake together with no exit and you get drift as a matter of arithmetic. Things come in faster than they go out. The pool of work expands while the pool of money and people does not. And because the new arrivals are assessed against an abstract bar rather than against the incumbents, the incumbents are never made to defend their place. They simply have it, by right of having had it yesterday.

The final ingredient is the most human: momentum gets mistaken for commitment. A piece of work that has been running for two years feels like something the organisation has decided to do. In truth, it may be something the organisation decided to start two years ago and has never once revisited. The decision was made under conditions — a strategy, a market, a set of assumptions — that have since changed beyond recognition. But the work carries on, propelled by its own history, and every month that it continues makes it feel more decided, not less.

1.2 What drift costs

It would be one thing if a drifted portfolio were merely untidy. It is not untidy. It is expensive, in ways that compound.

The first cost is dilution. A limited pool of money and skilled people, spread across too many things, is spread thin everywhere. The good initiatives are slowed because the people they need are also assigned to three other things; the weak initiatives are kept alive on a maintenance drip that produces nothing but keeps a claim on the budget. Nothing is properly resourced because everything is somewhat resourced. The organisation feels this as a permanent slowness — vast effort, little movement — and usually answers it by adding governance, which slows things further.

The second cost is the zombie. Every drifted portfolio contains initiatives that are neither alive nor dead: they have lost their sponsor, or their rationale, or their market, but no one has noticed or no one has acted. They consume budget, occupy people, generate status reports, and produce nothing anyone would miss. They are kept alive not by value but by the absence of a process for ending them. A portfolio can carry a surprising number of these before anyone notices, because each one is individually small and collectively invisible.

The third cost is strategic disconnection. Ask, of any large drifted portfolio, what proportion of its spend can be traced cleanly to a current strategic priority, and the

answer — if anyone is honest enough to compute it — is sobering. A great deal of the money is funding things that made sense under a previous strategy, or no strategy, or a strategy that was quietly abandoned. The link between what we say matters and what we actually fund has been severed, slowly, by accretion. The organisation has a strategy on a slide and a portfolio on a spreadsheet, and the two have not been on speaking terms for years.

The deepest cost of drift is invisible: it is the good work that was never funded because the money was already committed to work that had stopped deserving it.

This opportunity cost is the silent killer. While the organisation congratulates itself on being busy, genuinely higher-value investments — the ones that could move the strategy forward — never even make it onto the list because capacity is already locked up.

That cost of omission never appears on any report. Every euro tied up in an incumbent that no longer earns it is a euro unavailable for the new bet that might. The drifted portfolio is not merely carrying dead weight; it is crowding out its own future — declining next year's best work by default, not by judgement.

1.3 The tells

Drift announces itself, if you know the sounds. A few are so common they are almost diagnostic.

There is the initiative that only needs another ten percent to finish — and has needed exactly another ten percent for three consecutive years. The figure never changes, because it was never a measurement; it was the smallest number that sounded survivable.

There is the status report that has been green for eighteen months and turns red on the single day it is finally cancelled — proving, in one stroke, that the colour was never describing the work. It was describing the reporter's confidence that no one would look closely.

There is the strategic initiative that nobody can link to the strategy. Everyone agrees it is strategic. The word appears in its title. When asked which specific strategic objective it advances, the room becomes thoughtful, and then someone changes the subject.

There is the sponsor whose bonus depends on the initiative continuing, who is therefore the last person in the building able to assess, calmly, whether it should. And there is its cousin: the initiative too senior to question, sponsored by someone too important to be told no, which has become less a piece of work than a monument.

These are funny, in the bleak way that true things about working life are funny. But the humour is not the point. Each tell is a symptom of the same underlying disease: a portfolio in which the act of choosing has been replaced by the inertia of having chosen. The reports stay green not because lying is fun but because there is no safe mechanism for a piece of work to admit it has stopped being worth it. The ten percent never lands because no one is asking the only question that would expose it — would we start this

today, knowing what we now know?

1.4 Doing the right things

There is a comfortable belief, widespread and wrong, that the way to fix a struggling portfolio is to deliver better. Tighten the project management. Improve the estimates. Install a better tool. Run the standups more crisply. All of this is doing things right, and all of it is worth doing — inside an initiative that deserves to exist.

But no amount of execution excellence rescues a portfolio that is funding the wrong things. You can deliver a doomed initiative beautifully, on time and under budget, and at the end you will have a beautifully delivered thing nobody needed. The crispest standups in the world cannot make a zombie worth its budget. Precision in execution is wasted, and sometimes worse than wasted, when it is applied to work that should not be running at all — because it makes that work look healthy, which makes it harder to stop.

“You can execute the wrong portfolio flawlessly. The flawlessness is the trap.”

At the level of a single initiative, the question is are we doing this thing right? At the level of the portfolio, the question is entirely different: are these the right things to be doing at all? The two questions are not in tension, but they are not the same, and they are not solved by the same means. The portfolio's lever is not execution quality. It is choosing — and, far harder, re-choosing. The single most valuable act available to the person accountable for a portfolio is not making a running initiative go faster. It is recognising that a running initiative no longer earns its place, and moving its money to one that does.

That act is rare, and the reasons it is rare are the subject of much of this book. It is uncomfortable. It is political. It looks, to the untrained eye, like failure rather than discipline. But it is the entire game. A portfolio managed well is not one in which everything is delivered superbly. It is one in which, at every moment, the money is pointed at the best available set of things — and is moved, without sentiment, when the best available set changes.

1.5 The promise

If drift is the natural state, then the cure cannot be a one-off cleanup. Every organisation has, at some point, run the great rationalisation — the painful quarter where a task force reviews everything, kills a third of it, and declares the portfolio fixed. And every organisation has watched it drift straight back, because the forces that produced the sediment in the first place — additive intake, no exit, momentum mistaken for commitment — were never changed. They were merely interrupted. The cleanup treats the symptom. The drift resumes the following Tuesday.

The alternative is not a bigger cleanup. It is to stop treating the portfolio as a thing that gets decided and start treating it as a thing that gets continuously re-decided. A portfolio is not a plan you make and then execute; it is a standing question you answer, and keep answering: given everything we now know, and the limited money we have, what is the best set of things to be funding — and what should stop? Asked once, that question produces a clean-up. Asked on a rhythm, for ever, it produces a portfolio that cannot drift, because drift is exactly what regular re-choosing prevents.

That is the promise of this book. Not a better way to deliver, and not a one-time purge, but a way of running a portfolio as a living set of choices — surfaced, valued, funded, coordinated, and re-judged on a steady beat, with the money moving to follow the value wherever it goes. And, crucially, a way of doing this that is genuinely one method whatever the scale: the same discipline serves a single owner steering a handful of products and a global enterprise governing billions through tiered boards. What changes between them is the speed, the number of people in the room, and the weight of the paperwork. What does not change is the question.

The next chapter sets out that method — its single model, its unit of management, its two jobs, and the way one method bends to fit any scale. The rest of the book takes it stage by stage. But none of the machinery matters until the first idea is accepted: a portfolio you did not choose is one you are paying for, every day, in the work you cannot afford to start.

CHAPTER 2

The Canon

The Canon

2.1 A pool, not a list

The previous chapter ended with a promise: to run a portfolio not as a list that gets decided once and then defended, but as a living set of choices, re-made on a rhythm. This chapter sets out the language that makes that possible. It is the most definitional chapter in the book, and deliberately so. Everything that follows — every stage, every template, every argument about how to stop a piece of work — rests on a handful of terms used precisely. Get the terms right here and the rest of the book is the unfolding of a few simple ideas. Get them loose and the method dissolves into the usual fog.

Start with the founding idea, because everything else is a consequence of it.

A portfolio is a limited pool of resources. The job is to keep choosing the best mix of activities to fund for the best return — and then to make that chosen mix actually work together.

Every yes is a no to something else — whether or not the organisation ever names the alternative explicitly.

Two words in that sentence do the heavy lifting. Limited is the first. A portfolio is defined by scarcity; if money and people were infinite there would be nothing to manage, because you would simply do everything. The discipline exists only because the pool has a bottom. Keep is the second. The choosing is not an event. It is a standing activity. A portfolio is never finally decided; it is only ever continuously re-decided, because the world that justified yesterday's choices does not hold still.

This reframing — from list to pool, from decision to re-decision — is the whole shift. A list is a record of what was approved. A pool is a constraint you allocate, and keep re-allocating, against the best information you currently have. The list invites you to add. The pool forces you to choose.

2.2 The unit: the investment

To allocate a pool you need a unit — one comparable thing you put money into and expect a return from. In this method that unit is the investment, and the word is used with intent.

An investment is one thing the portfolio chooses to fund for a return. That is the entire definition, and its plainness is the point. It says nothing about size, duration, shape, or method of delivery. A two-week experiment is an investment. A four-year transformation programme is an investment. A standing product team funded quarter after quarter is an

investment. A piece of mandatory regulatory work is an investment — one whose "return" is the avoidance of a penalty, but an investment all the same.

This is why the book does not organise itself around projects, or programmes, or products, or initiatives. Those words describe how work is shaped and run, and they matter enormously to the people running it. But at the level of the pool they are noise. The portfolio does not care whether a thing is a project or a product; it cares what it costs, what it returns, how it fits the strategy, and what it risks. Programme, project, product and initiative are all simply types of investment.

"At the level of the pool, a programme and a two-week experiment are the same kind of object: money in, return expected, place to be earned."

Collapsing these terms into one is not pedantry. It is what makes a single method possible at every scale. The moment you maintain separate processes for "projects" and "products" and "programmes," you have three portfolios pretending to be one, three sets of rules, three queues that never have to compete with each other — and competition for the same money is exactly the thing a portfolio exists to stage. One unit, one queue, one set of questions.

2.3 The two jobs

A portfolio has two jobs, and confusing them is the source of a great deal of wasted effort.

The first job is to choose the right mix, and to keep re-choosing it. Out of everything that could be funded, which set actually should be — given the pool, the strategy, and what is now known? This is the job most people mean when they say "portfolio management," and it is the subject of roughly the first half of the method.

The second job is to make the chosen mix work together. Having chosen a set of investments, the portfolio must then coordinate between them: manage the dependencies that run from one to another, resolve the fights over the same scarce people, sequence work that cannot all happen at once, spot the duplication where two investments are quietly building the same thing, and clear the blockers that no single investment has the authority to clear alone. A perfectly chosen mix that is left to collide with itself will still fail.

Both jobs share one hard boundary, and it is worth stating once, plainly, because it governs every chapter that follows.

The portfolio coordinates between investments. It never runs the inside of any one investment. What happens within an investment — how the team plans, builds, and delivers — is delivery management, and it belongs to the investment, not the portfolio.

The line is between between and within. The portfolio's business is the relationships among investments and the allocation of the pool across them. The instant it reaches inside a single investment to direct how that investment does its work, it has stopped doing portfolio management and started doing the job of the people it funded.

This boundary is tested constantly. The most common failure of portfolio work is to quietly slip into running the investments instead of coordinating between them.

2.4 The Investment Loop

The method runs as a loop of six stages. They are continuous: each feeds the next, and the last feeds back into the first. The loop is the spine of the book, and Part II takes the stages one at a time. Here they are at a glance, in order, with nothing more than what each is for.

Stage	What it does
Direction	Turns strategy into the machinery of choice: the criteria, their weights, and the risk appetite that will govern every decision.
Demand	Surfaces and frames every candidate — and everything already running — into one comparable form, so the whole field of competitors for the money is visible at once.
Decide	Values the candidates and chooses the mix the pool can actually afford.
Fund	Releases the money on a rhythm, and funds uncertain work in stages so it can learn before it commits.
Steer	Makes the chosen mix work together while it runs — managing dependencies, contention, sequencing and blockers between investments.
Review	Re-judges the whole live set against everything now known, moves the money to where the value is, and stops what no longer earns its place.

The order matters, but the sequence is not the important property. The important property is that it is a loop, not a line. A line ends; a loop does not. Review does not close the portfolio down and send everyone home. It feeds straight back into Direction and Demand: the things it learns reshape the criteria, the money it frees re-enters the pool, the candidates it surfaces join the next round of choosing. There is no final state in which the portfolio is "done." There is only the next turn of the loop.

This is the structural answer to the drift described in the previous chapter. Drift happens when work enters a portfolio and never has to re-earn its place. The loop makes re-earning structural: every investment, running or proposed, comes back round to Review and Decide on a rhythm, and is judged again against everything else competing for the same pool. Nothing gets a permanent pass, because the loop keeps coming back around to ask.

A note on the word "Decide." As a stage it covers two distinct acts — working out what each candidate is worth, and then choosing which set the pool can afford. These are different enough that the book gives each its own chapter: How We Value, then How We Choose the Mix. When this book says the Decide stage, it means both together.

2.4.1 Which stage does which job

The two jobs map cleanly onto the loop, and seeing the map is the quickest way to hold the whole method in mind. Five of the six stages serve the first job — choosing the mix and re-choosing it. Direction sets the terms of the choice; Demand assembles the candidates; Decide values them and picks the set; Fund commits the money; Review re-opens the choice and moves the money on. One stage serves the second job. Steer, and Steer alone, is where the chosen mix is made to work together while it runs. This is why Steer reads differently from its neighbours: the others are about which investments; Steer is about the relationships between the investments already chosen. Both jobs run continuously, which is why a single stage can carry the whole of the second.

A caution against reading the loop too literally. Drawn as six boxes with arrows, it looks like a relay — Direction hands to Demand, Demand to Decide, and so on around the ring, once. It is not a relay. At any given moment a real portfolio has investments at every stage at once: some being framed, some being valued, some newly funded, many being steered, a few up for review. The loop describes the life each investment passes through and the rhythm on which the whole set is re-judged, not a queue the organisation marches through annually. The stages are always all running. What the rhythm governs is how often the portfolio stops, looks at everything together, and re-decides — which is the subject of the cadence dial and, ultimately, of the Review stage.

2.5 The two-way link with strategy

A portfolio does not invent its own purpose. It receives strategy — the organisation's view of what matters and why — and turns it into choices. That much is conventional. What is less conventional, and central to this method, is that the link runs in both directions.

Downward, strategy flows into the portfolio and sets the terms of choosing: what counts as valuable, which objectives weigh most, how much risk is tolerable. This is the work of the Direction stage.

Upward, the portfolio flows information back to strategy. A portfolio is the place where strategy meets reality and reality answers back. It is where you discover that the bet everyone was sure of returns nothing, that the unglamorous capability quietly underpins half the roadmap, that a market the strategy assumed is closing faster than anyone admitted. The portfolio is, among other things, the organisation's most honest instrument for learning what actually works — if it is allowed to report what it learns rather than what the strategy hoped.

"Strategy tells the portfolio what to value. The portfolio tells strategy what is true. A method that only listens downward is half a method."

A portfolio that only receives strategy and never feeds back becomes an order-taker, faithfully funding a plan long after the plan has been overtaken by events. The loop is two-way on purpose: Review's findings are not just inputs to the next funding round, they are signals to the strategy itself. How that feedback is captured and used is developed later, in the Review chapter and again in the operating model; what matters here is that it is part of the design, not an afterthought.

2.6 Two kinds of measures

The method keeps two kinds of measures rigorously apart, because collapsing them is one of the most common and most expensive confusions in portfolio life. They will be developed in full later; the distinction only needs to be planted here.

The first kind is selection criteria — the measures used to choose what to fund. Value, strategic fit, risk, cost: these are the lenses through which a candidate is judged worth funding or not. They are the basis of comparison across the whole pool, and they are set by Direction.

The second kind is performance measures — the measures used to check how a running investment is doing. Is it on track, on budget, hitting its milestones, delivering what it promised? These tell you about the health and progress of work already chosen.

The trap is to let the second quietly do the job of the first — to let "this project is going well" stand in for "this project deserves its funding." They are not the same question, and they do not connect directly.

Performance data does not decide funding on its own. It updates the value and cost numbers of a running investment. The funding decision then compares every investment — running and proposed — on the same terms.

That sentence is the whole relationship between the two, and it repays slow reading. A project that is performing beautifully may still lose its funding, because performing beautifully is not the same as being worth more than the alternatives now competing for the money. Good performance does not buy immunity; it simply updates the numbers that go into the next comparison. How performance feeds those numbers is the work of How We Value and Review; for now, hold only the wall between the two kinds of measures, and the single gate in it.

2.7 One method, three settings

The last piece of the canon is the one that lets a single method serve a single owner with a handful of products and a global enterprise governing billions. It is not three methods. It is one method with adjustable settings.

Think of the method as having a small set of dials. Four of them carry most of the weight: how often the portfolio re-decides (the cadence); how many people sit in the decision; how much formality and documentation the process demands; and how rigid the funding rules are. Turn those dials down and you get a fast, informal, lightly-governed portfolio. Turn them up and you get a slow, formal, heavily-governed one. The dials are continuous, but for everyday use they resolve into three named settings.

Dial	Lean	Managed	Enterprise
Cadence	Frequent and informal	A regular, scheduled beat	Periodic, tied to formal governance cycles
Who decides	One owner, or a small group	A standing board	Tiered boards, decisions nested by level
Formality	A page and a conversation	A short pack and a meeting	Full documentation and audit trail
Funding rigidity	Money moves freely as judgement changes	Money moves on the beat	Money moves through defined gates

The dials are not permanent. They can and should be adjusted deliberately as the organisation or its context changes.

The crucial property — and the reason this is one method and not three — is that the questions never change. At every setting you still surface demand, still value it against strategy-set criteria, still choose a mix against a limited pool, still coordinate the mix, still re-judge it and move the money. Lean does not skip the questions; it answers them faster, with fewer people, on a single page. Enterprise does not ask different questions; it answers the same ones more slowly, with more people, and more paper. A founder deciding over coffee which two of five product bets to fund this month is running the identical loop a multinational runs through its quarterly investment committee. The substance is the same. Only the speed, the number of people in the room, and the weight of the paperwork differ.

This matters in practice for two reasons. First, an organisation can run different parts of itself at different settings — a fast-moving product group on Lean while a regulated programme runs on Enterprise — without operating two incompatible systems, because underneath they are the same loop. Second, a portfolio can change its setting as it grows or as circumstances demand, dialling up formality when stakes and scrutiny rise and dialling it down when speed matters more, without throwing away the method and starting again. How to choose your own setting deliberately — rather than defaulting to whatever the last reorganisation left behind — is the work of the operating-model chapter near the end of the book. Here, only hold the shape of it: one method, four dials, three named settings, the same questions throughout.

2.8 What the rest of the book is

That is the canon. A portfolio is a limited pool, allocated and re-allocated as a set of investments, through a six-stage loop that runs two ways with strategy, doing two jobs — choosing the mix and making it work together — at one of three settings of a single method. Every chapter from here is the patient working-out of that sentence.

Part II takes the loop stage by stage, in order, and at each stage shows the Lean, Managed and Enterprise variations side by side rather than splitting the book by setting. It begins where the loop begins and where strategy first becomes machinery: with Direction — turning what the organisation says matters into the criteria, weights and appetite that will govern every choice the portfolio makes.

PART

Part II — Working the Investment Loop

CHAPTER 3

Direction

Direction

3.1 Where strategy becomes machinery

A strategy is a set of intentions. A portfolio is a set of choices. Between the two sits a translation problem that most organisations never solve: how to turn what the strategy says matters into something that can actually decide, month after month, which investment gets the money and which does not. Strategy is written in the language of ambition — grow in this market, modernise that platform, become the kind of company that does this. Choosing is done in the language of comparison — this one or that one, given that we cannot afford both. Direction is the stage that performs the translation. It takes the strategy as given and builds the apparatus that will apply it, automatically and consistently, to every decision the portfolio makes.

This is worth being clear about at the outset, because it marks the one thing Direction does not do.

Direction does not write the strategy. It assumes a strategy exists and turns it into the machinery of choice: the criteria, the weights, the risk appetite, and the funding buckets. If there is no strategy, that is a problem to solve before the portfolio can function — not a gap the portfolio can fill by inventing direction of its own.

The output of Direction is not a plan, a roadmap, or a list of approved work. It is a small set of durable instruments that the rest of the loop runs on. Get them right and the later stages have something solid to push against. Get them wrong — vague criteria, weights nobody believes, a risk appetite that permits everything — and every subsequent decision becomes an argument, because there is no agreed basis on which to settle it. Most portfolio dysfunction that looks like a problem of deciding is really a problem of Direction: the decisions are hard because the basis for them was never built.

There are four instruments. We take them in turn.

3.2 The reference framework

The first and most important instrument is the set of selection criteria — the lenses through which every candidate investment is judged. In the canon these were named as a kind: the measures used to choose what to fund, as distinct from the performance measures used to check how chosen work is doing. Here we make them concrete.

Four families of criterion cover almost every portfolio. They are not a rigid template — an organisation may name them differently or split them finer — but if your framework is missing one of these, it is probably missing something it will later regret.

Criterion	The question it asks
Value	How much benefit does this create, of the kind the organisation counts — and how confident are we in that number?
Strategic fit	How directly does this advance a current strategic priority, as opposed to merely being useful?
Risk	How likely is this to fail, and how damaging if it does — to the investment itself and to the rest of the portfolio?
Cost	What does this consume from the limited pool — money, yes, but also scarce skills and management attention?

The agreed set of criteria, together with the weights attached to them, is what this book calls the reference framework. It is the constitution of the portfolio. Every investment is judged against it; every funding argument is, ultimately, an argument conducted in its terms. Because it is constitutional, it should be small, stable, and understood by everyone who proposes or decides. A framework with fourteen criteria is not more rigorous than one with four; it is less, because no one can hold fourteen things in mind while comparing two investments, so in practice they collapse it back to the three or four that actually move the decision — only now they do so privately and inconsistently. Keep the set small enough to be used honestly.

“A criterion nobody can apply consistently is not a criterion. It is a decoration that makes the spreadsheet look thorough.”

Two cautions on building the framework. The first is the vanity criterion — the lens that sounds impressive and measures nothing, like “innovativeness” or “transformational impact,” which in practice means whatever the most senior person in the room wants it to mean. If a criterion cannot be applied to two real investments to produce a defensible difference between them, it is not earning its place. The second is the double count — value and strategic fit, in particular, tend to bleed into each other, so that a strategically aligned investment scores well twice for the same reason. Keep the criteria genuinely distinct, each asking a question the others do not.

What this chapter deliberately does not do is explain how to score an investment against these criteria, or how to combine the scores into a ranking. That is the work of valuation, and it has its own chapter. Direction's job is to agree what will be measured and how much each matters — not to do the measuring.

3.3 Strategy sets the weights, not the scorecard

Here is the mechanism that makes the framework strategic, and it is the most important idea in the chapter.

The criteria are stable. Value, strategic fit, risk and cost matter to every portfolio, in every year, under every strategy. What changes with strategy is not which criteria you use but how much each one weighs. A strategy that prizes growth above all will weight value and strategic fit heavily and tolerate more risk. A strategy of consolidation after a hard year will weight cost and risk heavily and demand near-certain value. The same four lenses, in radically different proportions, produce radically different portfolios — which is exactly right, because the strategy is different.

Strategy sets the weights, not the scorecard. The criteria are the constant; the weights are the variable. When the strategy changes, you do not rebuild the framework — you re-balance it.

This separation does a great deal of quiet work. It means the framework is durable — you are not redesigning your decision apparatus every time the strategy is refreshed, only re-tuning it — which keeps the portfolio comparable across time and stops every strategy refresh from triggering a methodological civil war. It means the strategic intent is expressed in one visible, debatable place: the weights. And it means a change of strategy can be implemented in the portfolio in an afternoon, by re-weighting, rather than over a year of trying to argue individual investments into or out of favour.

Consider what re-weighting actually does. Suppose strategic fit carried a weight of twenty per cent, and a strategy refresh elevates a particular market to the centre of the company's plans. Raise strategic fit to thirty-five per cent and several things happen at once, automatically, the next time the portfolio is judged: investments serving that market rise in the ranking, investments that were coasting on raw financial return without strategic relevance fall, and the marginal cases — the ones near the funding line — re-sort themselves. No one had to relitigate each investment. The weight change did it. That is what it means for strategy to drive the portfolio: not exhortation, not a stirring all-hands, but a number that changes what wins.

How often should the weights move? Less often than people fear and more often than most organisations manage. Weights should be reviewed when the strategy genuinely changes — and only then. Re-weighting every quarter because the wind shifted is its own pathology; it makes the portfolio unstable and teaches everyone that the framework is theatre. But a strategy that has visibly moved while the weights sat still is the surest sign that the reference framework has quietly detached from reality, and that the portfolio is now optimising for a strategy the company no longer has.

3.4 Risk appetite

The third instrument is the portfolio's risk appetite — a deliberate statement of how much risk, and what kinds, the portfolio is willing to carry in pursuit of return.

Risk already appears as a selection criterion, where it shrinks the attractiveness of an individual investment. Risk appetite is a different and higher-level thing. It is not about any single investment; it is about the shape of the whole. A portfolio of nothing but safe, certain, incremental investments is not a prudent portfolio — it is a portfolio quietly

guaranteeing that nothing important will ever happen, because everything important is uncertain at the outset. A portfolio of nothing but bold bets is not a brave portfolio — it is one likely to deliver nothing at all. Appetite is the statement of where, between those poles, this portfolio is meant to sit.

A useful risk appetite does two things. It permits — it tells the portfolio that a certain amount of failure is not only acceptable but expected, which is the only thing that makes genuine bets fundable. And it forbids — it draws the lines the portfolio will not cross, whatever the apparent return: the concentration it will not allow, the categories of risk it will not accept, the floor of run-the-business reliability it will not trade away for upside.

“A risk appetite that permits nothing funds only the obvious. A risk appetite that forbids nothing is not an appetite — it is a shrug with a letterhead.”

The commonest failure here is the appetite statement that has been carefully drafted to mean nothing. “We have a balanced appetite for risk” is the corporate equivalent of saying the weather is sometimes nice. A real appetite is specific enough to change a decision: it says, in effect, we will tolerate this much exposure to this kind of failure, and no more, and here is the room we are deliberately leaving for bets that may not pay off. When the time comes to choose the mix, the appetite is what keeps the portfolio from quietly defaulting to a wall of safe incremental work — the natural gravitational pull of any group of people whose mistakes are punished more reliably than their successes are rewarded.

How risk appetite is applied — how it shapes the actual balance of the chosen set — belongs to the chapter on choosing the mix. Direction's job is to set it: to make it explicit, specific, and agreed, before any individual investment is on the table to distort the conversation.

3.5 Funding buckets

The fourth instrument is optional, and powerful, and frequently abused. It is the funding bucket.

A funding bucket is a deliberate pre-commitment that slices the total pool into parts before any investment competes for it. The classic split is three ways: run the business (keeping the lights on, the existing operation funded and reliable), change the business (improving and extending what exists), and innovation (bets on what might come next). Other splits exist — by market, by horizon, by capability — but the principle is constant: you decide, in advance and at the level of the whole portfolio, roughly how much of the pool each part should get, rather than letting every kind of work fight every other kind in a single undifferentiated contest.

The reason to do this is protective, and it solves a specific, predictable failure. In a single undifferentiated pool, the safe, near-term, easily-quantified investment beats the

uncertain, long-term, hard-to-quantify one almost every time — not because it is more valuable, but because its value is easier to defend in the room. Run-the-business work, with its concrete and immediate justification, will reliably starve innovation, which trades in possibility. Left alone, the pool eats its own future, one defensible quarter at a time. Buckets prevent this by ring-fencing a share of the money for the kinds of work that would otherwise always lose the argument.

The discipline that makes buckets work — and the discipline that most organisations abandon — is captured in a single rule.

Investments compete within a bucket each period. Bucket sizes are resized between periods. The buckets protect strategic and uncertain bets from being crowded out — but they are not permanent fiefdoms, and they are not an escape from re-allocation.

Both halves of that rule are load-bearing. Compete within means a bucket is not a guarantee of a quiet life: the innovation bucket's contents still fight each other for the innovation money, on the same criteria as everything else, and weak bets in a protected bucket are still weak bets. The protection is for the category, not for any investment inside it. Resize between means the bucket boundaries are themselves subject to the strategy and the weights: a strategy that shifts toward growth should grow the change and innovation buckets at the expense of run; a year that demands consolidation should do the reverse. A bucket whose size has not changed in five years is no longer a strategic instrument. It is a baron's territory — and the surest sign of a portfolio where the buckets have stopped serving the strategy and started serving whoever owns them.

Buckets are not free. They add a layer of structure, and a small portfolio may not need them at all — when there are eight investments and one decision-maker, the protective function can be held in a single head. The question of whether to use buckets, and how the within-competition and between-resizing actually run, connects directly to choosing the mix, and is developed there. Direction's contribution is the framework: deciding whether to slice the pool, on what lines, and roughly in what proportions, as an expression of the strategy.

3.6 Direction is where the loop closes

Direction is drawn as the first stage, and in the order of explanation it is. But it is also the last. The loop, as the canon insisted, is not a line: Review feeds back into Direction, and Direction is where that feedback has to land or be lost. This is the second, less obvious reason the weights move — and it is the one organisations almost always miss.

Weights move on a change of strategy. That much we have covered: a deliberate, top-down re-tuning when the company's intent shifts. But weights should also move on evidence — on what the portfolio has learned about what actually works. The two-way link with strategy is not a slogan; it is a mechanism, and Direction is the mechanism's far end. When Review reveals that a whole class of investment the framework rated highly has reliably failed to deliver the value it promised, that is not merely a fact about those investments. It is a fact about the framework. A criterion that keeps selecting losers is

mis-weighted, or mis-defined, or measuring the wrong thing — and Direction is where it gets corrected.

The reference framework is not just an input to the portfolio. It is the portfolio's memory of what it has learned about value. If nothing the organisation discovers ever changes the criteria or the weights, the loop is not closing — and the portfolio is learning nothing.

This sets up a useful tension, and Direction's craft is holding it. The framework must be stable enough to keep the portfolio comparable across time and to resist every passing enthusiasm — and alive enough to absorb hard evidence about what the organisation's bets actually return. Too stable and it fossilises, faithfully selecting the same kind of work that has been quietly disappointing for years. Too reactive and it becomes noise, re-weighted on the strength of the last loud success or failure, with no memory at all. The standard to hold is simple to state and hard to live: change the framework when the strategy changes, or when the evidence is strong enough to overturn a belief the framework encodes — and not otherwise.

3.7 Direction at three settings

Everything above holds at every scale. What changes between Lean, Managed and Enterprise is not the instruments but their weight — how many criteria, how formally the weights are set, how elaborate the appetite, whether buckets exist at all.

Instr ume nt	Lean	Managed	Enterprise
Criter ia	Three or four, written on a single page	A defined framework of four to six, documented	A full framework, formally governed, with definitions and scoring guidance
Weig hts	Set by the owner, in their head or on the page	Set by the portfolio board and recorded	Set with the strategy function, ratified by the governing board, version-controlled
Risk a ppetit e	A sentence or two the owner actually believes	A short statement, reviewed each period	A formal appetite framework with explicit limits and categories
Bucke ts	Often none, or an informal split held in mind	Two or three buckets, sizes reviewed each period	Multiple buckets, sometimes nested by tier, formally resized on the cycle

The temptation at the Lean end is to skip Direction entirely — to start judging investments before agreeing what matters. This feels efficient and is not. A single owner with four products and no written framework is still applying criteria and weights; they are simply applying them invisibly, which means inconsistently, and which means they cannot explain or defend a choice when challenged. The Lean version of Direction is small — a few criteria, a sentence of appetite, no buckets — but it is not absent. It fits on one

page, and that page is worth more than its size suggests.

The temptation at the Enterprise end is the opposite: to mistake elaboration for rigour, building a forty-criterion framework with a scoring manual that no one reads and a risk appetite statement long enough to require its own table of contents. The discipline at scale is to keep the constitutional part — the criteria and weights that actually decide — small and sharp, and to let the formality live in the governance and documentation around it, not in the framework itself. A bigger organisation needs more process around its Direction. It does not need more criteria.

With Direction set — the criteria agreed, the weights tuned to the strategy, the appetite stated, the buckets drawn — the portfolio has a basis on which to judge. What it does not yet have is anything to judge. That is the next stage. Demand surfaces the candidates and frames them into a form the reference framework can be applied to — including, crucially, the work already running, which must re-enter the contest rather than assume its place in it.

CHAPTER 4

Demand

Demand

4.1 The visibility problem

Direction has built the apparatus: the criteria, the weights, the appetite, the buckets. It cannot decide anything yet, because there is nothing in front of it. The portfolio needs candidates — things to judge, compare, and potentially fund. Demand is the stage that surfaces them. Its job sounds administrative and is not. It is, in practice, one of the most consequential acts in portfolio management, because what does not enter the demand stage does not get chosen, and what does not get chosen does not get funded. The shape of the demand that the portfolio sees determines the shape of the portfolio it can build.

The visibility problem is this: in any real organisation, the full population of things competing for the money is never visible in one place. Some candidates are formally proposed and awaiting review. Some are in planning, not yet named a candidate but already consuming preparation time. Some are already funded and running, competing for the same pool just as surely as any new arrival — but because they started, they have stopped being treated as competitors. Some are the things nobody formally proposed because the person who had the idea knew it would lose the argument with the current incumbents, and so decided not to bother. Some are "already decided" — investments that a senior sponsor committed to in a hallway and which now arrive in the portfolio demanding only the formality of being logged.

Demand's job is to surface all of them — the formal, the informal, the in-flight, and the merely expected — into a single, visible pool. Not so they can all be funded. So they can all be judged.

Nothing competes off the books. Every investment — proposed, in-flight, or already running — must re-enter the demand pool and take its place in the comparison. An investment that assumes its place is an investment that has escaped the portfolio's discipline.

This is uncomfortable in practice, and it should be. The running investment that is brought back into comparison is no longer simply "ongoing work." It is a candidate, re-entering the contest, required to justify its continued claim on the pool against everything else that wants the money. That discomfort is the point. A demand stage that only processes new proposals, and leaves the incumbents untouched, is not running a portfolio. It is running an intake process alongside a portfolio that nobody manages.

4.2 What makes a candidate an investment

Not everything that arrives at the demand stage is ready to be judged. Some are no more than an idea and a sponsor name. Some are a PowerPoint assembled by someone who wants the money enough to get a meeting but has not yet done the work of making the case comparable. To enter the comparison, a candidate must state a minimum: the information without which the reference framework cannot be applied.

That minimum has four elements, and they map directly onto the four families of selection criterion.

Element	The question it answers
What and why	What is this investment, and what return does it claim to create — for whom, and measured how?
Strategic fit	Which strategy objective does it advance, and how directly?
Cost and resource	What does it require from the pool — money, time, people, and any other scarce resource it competes for?
Risk and confidence	How uncertain is the return, and what is the main thing that could make it fail?

These four elements constitute the Investment One-Pager — the standard framing document that every investment must produce before it is eligible for comparison. The name is intentional. If the case cannot be stated on one page, it has not been thought through enough to be compared to others that have been. A candidate that requires forty slides to make its argument is a candidate that has confused advocacy with analysis. The One-Pager is not a business case; it is a comparable entry. The full case is developed if and when the investment is selected.

"The purpose of the One-Pager is not to win the argument. It is to enter the comparison on the same terms as everything else."

One element of the minimum deserves specific attention: the return claim must specify what value will be created, for whom, and how the organisation will know it arrived. "Improve customer satisfaction" is not a return claim. It is an intention. "Reduce average complaint resolution time by thirty per cent, measured on the existing support system, by the end of the next financial year" is a return claim — one that can be assessed for plausibility and tracked after delivery. Candidates that cannot state their return in terms that can later be verified are, at minimum, not yet ready to be compared, and may be trying to fund something that nobody will have to account for.

4.3 In-flight work enters the same door

Running investments produce the most important — and least comfortable — version of the demand question. An in-flight investment is already consuming the pool. It already has a team, a sponsor, a status colour, and a set of expectations attached to it. The political weight of all of that means that, in almost every portfolio that lacks a formal re-entry discipline, it never actually re-enters the comparison. It is assumed to continue until it finishes, or until something dramatic enough to force a decision makes the conversation unavoidable.

This assumption is the mechanism that keeps zombie investments alive. Nothing decided to keep funding them; no one decided to stop. The money just kept flowing because no one made the question visible.

The method's answer is direct: in-flight investments complete a One-Pager too, updated to reflect current knowledge — what has been learned since they started, what the value still to come actually looks like now, and what the cost to complete still is. They enter the demand pool on the same terms as every new candidate. They do not get a separate queue, a lower bar, or the benefit of the doubt that came with their original approval. The money already spent on them gets no vote; only the value still to come, against the cost still to spend, is the relevant number.

This is where Principle 4 — every activity must keep earning its place, judged on value to come versus money to spend, not money already spent — moves from a statement to a practice. Demand is the stage that makes it structural.

4.4 The intake log and the single front door

Demand works only if candidates must enter through a single point. An organisation in which investments can be approved via the portfolio process, or via a direct executive decision, or via a budget exception, or via "we already started so it's too late to say no" has not got a portfolio — it has got a portfolio process and several ways around it, and the ways around it will be used whenever the portfolio process looks likely to say no.

The intake and demand log is the instrument that enforces the single front door. It is a running record of every candidate — new and in-flight — that has been formally surfaced for comparison, at what point in its development, and what happened to it. It is not a project register. It is a visibility tool: it shows the portfolio owner or board, at any point, the full population of what is competing for the money, what stage of framing each is at, and what is pending a decision.

The demand log is not bureaucracy. It is the mechanism that makes "nothing competes off the books" real rather than aspirational. An investment that is not on the log is not in the portfolio process — and should be challenged the moment it becomes visible.

The log also prevents a quiet failure mode: the candidate that gets as far as the door, is informally discouraged, and disappears without ever being formally declined. Those invisible rejections matter, both because the person who proposed the investment deserves a clear answer, and because the portfolio loses the signal of what was never surfaced — which, over time, tells you something important about where the real demand

lies and what the portfolio process has trained people not to propose.

4.5 Flagging uncertain bets

Not every investment arriving at the demand stage is a straightforward claim about a known return. Some are bets — investments whose return depends on things nobody currently knows: whether the market will respond, whether the technology will work at scale, whether the underlying assumption the strategy is built on is actually true. These are not inferior investments. In many portfolios they are the most strategically important ones — the ones that change what is possible rather than incrementally improving what exists. But they need to be handled differently from the outset, and the demand stage is where the difference is first named.

At intake, an uncertain bet is flagged as a candidate for fund-to-learn treatment. This is not a valuation; it is a classification. The flag says: this investment's return cannot be reliably estimated in the normal way, because it depends on information that does not yet exist and can only be created by doing some of the work. What the portfolio needs to know about this candidate, before it can be judged in the normal way, is what the next increment of learning is worth and how much it costs to acquire. That is a question for valuation — for the How We Value chapter — and the flag sends the investment to that process on the right terms, rather than forcing it through a financial appraisal model that will produce a confident number based entirely on assumption.

The practical effect of the flag is modest at the demand stage: the investment is admitted to the pool, noted as uncertain, and queued for the right kind of valuation. The deeper effect is cultural. An organisation that allows uncertain bets to enter the demand process as a legitimate category — rather than requiring them to pretend to a certainty they cannot have — is one that will fund the kind of work that might actually change its future. An organisation that runs every candidate through the same financial appraisal model, regardless of how uncertain its return is, will reliably fund the certain-return incumbents and defund everything that looks like a bet, because the bet will always produce a worse-looking number when forced into a formula designed for a different kind of investment.

4.6 Demand at three settings

The question of how much process to put around demand is the one most organisations get wrong in both directions at once: too little discipline at the Lean end, where candidates surface informally and the intake log lives in someone's memory; too much paperwork at the Enterprise end, where the One-Pager has become a seventy-five-slide template that takes three months to produce and teaches sponsors to avoid the process entirely.

Aspect	Lean	Managed	Enterprise
Front door	A direct conversation with the owner, or a short shared list	A defined intake process with a standard form and a submission window	A formal intake gate, governed, with submission guidance and a review calendar
One-Pager	Genuinely one page, verbal if the organisation is small enough	A short structured template covering the four elements	A fuller template with supporting analysis, reviewed before it reaches the board
Intake log	A shared list the owner maintains	A structured log, visible to all who propose and decide	A governed register with version history, updated each cycle
In-flight re-entry	Periodic conversation: is this still the best use of the money?	A formal update against the standard template each period	A full re-entry submission on the standard cadence, treated identically to new proposals
Uncertain-bet flag	Owner's judgement at intake	A tick-box on the template, reviewed at the intake gate	A formal classification with a specific fund-to-learn pathway

At the Lean end, the most important discipline is the in-flight re-entry. It is the easiest to skip — the owner knows what is running, thinks they are on top of it, and the formal comparison feels unnecessary overhead. It is, in practice, the single most valuable habit a small portfolio can build: the regular moment of asking whether the current set, visible all at once, is still the best use of the available money. Without it, the Lean portfolio drifts exactly as its larger cousins do, just faster.

At the Enterprise end, the challenge is the opposite. When the intake process becomes onerous enough that sponsors prefer to route funding through other channels, the demand stage has achieved the reverse of its purpose: it has made the portfolio's visibility worse rather than better. The test is simple — if proposals are routinely arriving outside the intake log, the intake process is too hard.

With the demand pool visible — the new candidates framed, the in-flight work re-entered, the uncertain bets flagged — the portfolio can move to the next question: what is each of these actually worth, on comparable terms? That is the work of the following chapter.

CHAPTER 5

How We Value

How We Value

5.1 The comparison problem

Demand has assembled the pool: new candidates framed on the One-Pager, running investments re-entered, uncertain bets flagged. The pool is now visible. It is not yet useful, because the investments in it are not yet comparable. A data platform modernisation, a new market entry bet, a mandatory compliance upgrade, and a cost-reduction programme are all described in the language of their own logic, on their own terms, with their own units of value. The portfolio cannot choose between them until they can be placed on the same scale.

Valuation is the act of translating every investment into comparable terms, using the criteria and weights the reference framework set in Direction. This chapter sets out how to do that – which models are available, when to use each, what forward value means in practice, and how to value the investments whose return genuinely cannot be known in advance. It ends with the numbers the next stage needs: a comparable value for every candidate in the pool, arrived at honestly.

What this chapter deliberately does not produce is a ranked list or a funding decision. Those are the work of the following chapter. Valuation's output is the inputs to a choice, not the choice itself.

5.2 The models and when to use them

There is no single correct way to value an investment. There are several models, each well-suited to a different kind of investment and a different kind of organisation. The mistake is to use one model for everything – to push every investment through a financial appraisal because the finance team is comfortable with spreadsheets, or to use weighted scoring for everything because it is easy to explain. The skill is knowing which model to reach for, and why.

Four families of model cover the portfolio's needs.

5.2.1 Weighted scoring against the reference framework

The most direct application of Direction's work. Each criterion in the reference framework is scored, the scores are multiplied by the strategy-set weights, and the totals become the basis of comparison. The appeal is transparency: the scores are visible, the weights are agreed and published, and the outcome of changing a weight is immediately observable.

Weighted scoring works well when the investments being compared are genuinely heterogeneous – different in shape, size, sector, and type – and where the portfolio

needs a single comparable number that encodes strategy, not just financial return. It is the default for most portfolios because it can be applied to anything the reference framework can assess.

Its weakest points are that scores for criteria like "strategic fit" or "value" depend on consistent judgement across scorers — a criterion that means different things to different assessors produces noise, not signal — and that it can become a mechanical ritual that produces whatever number was wanted before the scoring started, if the review process around it is weak. The model is only as rigorous as the conversation that produced the weights and validates the scores.

5.2.2 Cost of Delay and WSJF

Cost of Delay is the answer to a question that financial appraisal rarely asks: what is the cost to the organisation of not doing this investment now? Every day an investment is not running is a day its value is not being created. That delay has a cost — in revenue foregone, in competitive position not captured, in user problems not solved — and that cost can be expressed as a rate: the value lost per unit of time the investment is deferred.

Weighted Shortest Job First (WSJF) turns that insight into a ranking mechanism: divide the Cost of Delay by the duration or effort required to deliver the investment, and the result is a measure of the rate at which value is created relative to the time the delivery takes. A small investment with a large Cost of Delay and a short delivery duration should, in most cases, be done before a large investment with a moderate Cost of Delay and a long delivery duration — even if the large investment's total value is greater — because the small one creates its value faster and frees the resource sooner.

Cost of Delay makes time visible. A portfolio that never asks "what does delay cost us?" is implicitly treating delay as free — which it is not, and which leads to queues of high-value work waiting behind large, slow incumbents.

Cost of Delay and WSJF are most powerful in portfolios with a high rate of demand — technology and product portfolios, especially — where the sequencing of work and the time-to-value of individual items is a significant driver of outcome. They are less useful where most investments are long-running, structurally similar, and where the comparison is across investment categories rather than within a queue of comparable work items.

5.2.3 Value-versus-effort (the simple portfolio map)

A plotting of investments on two axes — value (the return expected) against effort or cost (the resource required) — is the most accessible of the valuation tools and one of the most useful as a portfolio-level view. High value, low effort: do these first. Low value, high effort: these are the candidates for stopping. The middle is where most of the interesting decisions live.

The value-versus-effort map does not pretend to precision. Its power is transparency at the portfolio level — it makes the shape of the whole set visible in a way that numbers in a spreadsheet do not. Used well, it is a conversation starter that reveals outliers, clusters,

and imbalances that nobody had previously named. Used badly, it is a scatter chart produced in a workshop and ignored by Tuesday.

It works best as a complement to a more rigorous valuation for individual investments, and as the view the governance board uses to interrogate the whole set rather than the basis on which individual investments are funded or stopped.

5.2.4 Financial appraisal — ROI, NPV, payback

The traditional toolkit of investment finance. Return on investment expresses the expected return as a proportion of the cost. Net present value discounts expected future cash flows to their present value, accounting for the time-value of money and the organisation's cost of capital. Payback period states how long before the investment recovers its cost.

These models are powerful when the value of an investment is primarily financial and can be estimated with reasonable confidence — when the investment will generate or save measurable money, in a timeframe that can be modelled, against assumptions that can be defended. They are the right tool for investments where finance has the data to run them honestly.

They become a liability when they are required everywhere, including where the value is not primarily financial (a culture change programme, a research experiment, a strategic capability) or where the estimates cannot be honest (a five-year NPV built on three layers of assumption, each plausible and collectively fantasy). A confident number produced by a bad model is worse than an honest acknowledgement of uncertainty, because the confident number goes into a spreadsheet and gets compared on equal terms with numbers that actually mean something.

“Financial appraisal is the right tool for investments whose value is financial and estimable. Required everywhere else, it teaches people to produce the number that justifies the investment they have already decided to make.”

5.3 Which model for which setting

The choice of model should match the investment and the organisation's capability to use it honestly.

Model	Best for	Weakest when
Weighted scoring	Heterogeneous portfolios; where strategy alignment matters as much as financial return	Scorers apply criteria inconsistently; weights are gamed

Model	Best for	Weakest when
Cost of Delay / WSJF	High-flow product and technology portfolios; sequencing within a queue	Investments are long, heterogeneous, or not sequencing-dependent
Value-versus-effort	Portfolio-level view; conversation with the board; identifying outliers	Used as the primary valuation for individual funding decisions
Financial appraisal	Investments with measurable financial returns and defensible estimates	Value is not financial; estimates are multi-layered assumption

In practice, most portfolios use weighted scoring as the primary comparison model and supplement it with Cost of Delay for sequencing decisions within a category, financial appraisal where the numbers are genuinely available, and the value-versus-effort map as the governance view of the whole set. The principle is not model purity but model honesty: use the approach that best represents what this investment's return actually looks like, and resist the institutional pressure to use a single model for everything because it makes the spreadsheet tidy.

5.4 Applying the strategy-set weights

Whichever model is primary, the strategy-set weights from the reference framework must be applied. This is the mechanism, established in Direction, by which the organisation's strategic intent enters the valuation. Without it, valuation is optimising for the wrong thing – typically, the easiest-to-defend financial return – regardless of what the strategy says should matter.

Applying the weights is not complicated, but it requires discipline. The score for each criterion is multiplied by the criterion's weight before the totals are summed. A criterion weighted at thirty-five per cent has more than twice the influence on the outcome of a criterion weighted at fifteen per cent. Changing the weights changes who wins – which is exactly what was established in Direction, and is the point. The person running the valuation does not adjust the weights to get the desired outcome; that conversation happened, under different conditions, when the weights were set.

The integrity of the comparison depends on the weights being applied consistently across all investments in the pool. An investment that gets its strategic fit scored differently from its competitors because its sponsor is more articulate, or because the scoring session for it was chaired by someone with a different view, is not competing on the same terms. Consistent application is a process discipline, and it is one the intake log and the comparison meeting must enforce.

5.5 Forward value: the only number that matters for the decision

Whatever model is used, the number it produces must be a number about the future, not the past.

Forward value is the value still to come, measured against the cost still to spend. It is the only relevant quantity for a funding decision, because a funding decision is about whether to continue committing the organisation's limited pool to this investment going forward. The money already spent is gone. It cannot be recovered by continuing, and it cannot be avoided by stopping. It gets no vote.

This is one of the most important ideas in the method, and one of the most routinely violated in practice. The violation looks like this: an investment has consumed eight million over three years. It has delivered little of what it promised. The forecast to complete has grown twice. In the valuation review it is nonetheless compared against candidates with much lower total investment, because "we've already spent eight million — we can't walk away now." The eight million is not evidence of the investment's future value. It is evidence of its history. Forward value says: set the eight million aside. What is this investment worth, from today, against what it will cost, from today? If that number is worse than the alternatives in the pool, the investment does not earn its place — regardless of what has already been spent on it.

The money already spent is gone and gets no vote. Forward value — value still to come versus cost still to spend — is the only basis for a funding decision. Sunk cost is a fact about the past, not a claim on the future.

Forward value for a running investment is not the original business case. It is an updated estimate of the remaining benefit, based on what has actually been learned about the investment's performance, against the remaining cost to complete. The Demand stage structured the update; the valuation models translate it into a comparable number. The How We Value stage is where sunk cost thinking is confronted and, with discipline, overcome.

5.6 Valuing uncertain bets as options

The investments flagged at the demand stage as uncertain bets cannot be honestly valued by the standard models, because their return depends on information that does not yet exist. Forcing them through a financial appraisal or a weighted score produces a number, but not a meaningful one: it encodes assumptions about a future that the bet itself is designed to test.

The right conceptual frame is option value. An uncertain investment is not a commitment to a known return; it is the purchase of the right to learn something, and then to make a better-informed decision on the basis of what was learned. The question is not "what is this investment worth if it works?" — that question cannot be answered honestly. The question is: "what is the next increment of learning worth, and what does it cost to acquire?"

In practice this means valuing uncertain bets in two parts. First, the small amount of money and time required to do the next piece of learning — the prototype, the experiment, the market test, the technical spike. Second, the option that learning creates: if the learning is positive, the right to commit to the next stage; if negative, the ability to

stop having spent the minimum. The value of the option is not the full projected return of the investment if it goes well; it is the information value of knowing which it is likely to be.

“Funding an uncertain bet to learn is not a commitment to deliver. It is the purchase of better information. The decision to deliver comes after the learning, not before it.”

This framing resolves a common governance problem. A governance process that treats every investment as a commitment to a deliverable, and therefore demands a full business case, will never fund uncertain bets honestly — because an honest uncertain bet cannot produce a credible full business case. The organisation ends up either declining the bet (and missing what might have come next) or funding it on the basis of fabricated confidence (and being surprised when reality differs from the case). Valuing bets as options creates a third path: fund the learning, genuinely, without demanding the pretence of certainty, and reserve the larger commitment for after the information is in.

5.7 How performance data updates the value

Running investments are valued forward, as established above. But the forward value is not simply what the original case said it would be, updated for time. It is updated by what has actually happened: whether the value being created is tracking as expected, whether costs have changed, whether the assumptions the case depended on have held.

Performance data — the health and progress of a running investment — enters the valuation through a specific gate. It does not, by itself, change the funding decision. It updates the forward value and forward cost of the investment. Once updated, that investment re-enters the comparison on the same terms as every other candidate in the pool, with its revised numbers. The comparison, using the same reference framework and weights, then determines whether it still earns its place.

This is the mechanism that was named at a glance in the canon, now made operational. A project that is performing well — on time, on budget, delivering as promised — will have its forward value confirmed and its forward cost unchanged: good performance is reflected in the numbers, and the numbers go back into the comparison. A project that is performing badly will have its forward value downgraded and its forward cost increased: bad performance is also reflected in the numbers, and the numbers go back into the comparison. Neither "performing well" nor "performing badly" by itself decides the funding; the comparison decides it, using the updated numbers as inputs.

5.8 How We Value at three settings

Aspect	Lean	Managed	Enterprise
Primary model	Weighted scoring, lightweight; or owner's structured judgement	Weighted scoring with defined criteria and a scoring session	Weighted scoring as the standard; supplemented by financial appraisal where data supports it
Forward value	A pragmatic estimate: what is realistically left to gain, what does it cost to get there	A documented forward value calculation for each investment	A formally updated estimate each period, reviewed by the finance function
Uncertain bets	Owner decides the learning investment and the stop criteria	Flagged at intake, valued as a learning tranche with an explicit decision point	A governed fund-to-learn pathway with tranche approvals and learning milestones
Performance update	A conversation about whether the value claim still holds	A structured update on the standard template each period	A formal reforecast reviewed before each funding cycle

The temptation at the Lean setting is to skip the structured comparison and rely on the owner's intuition about what is most valuable. Intuition informed by the right criteria and honest forward value is excellent. Intuition uninformed by those things is habit, applied with more confidence than it deserves. Even a simple, quick scoring session against the reference framework does more than it might appear to do: it forces the criteria to be named, the weights to be applied, and the comparison to be explicit — which means it can be challenged and improved, rather than reconstructed after the fact to justify a decision already made.

With each investment in the pool carrying a comparable value — and each uncertain bet identified and sized for what its first learning tranche is worth — the portfolio can now do the hardest thing: choose. That is the subject of the following chapter.

CHAPTER 6

How We Choose the Mix

How We Choose the Mix

6.1 From values to a portfolio

Valuation has given every investment in the pool a comparable number — a forward value, arrived at honestly, using the model suited to what each investment actually is. The pool is now ready for the hardest question: which of these should actually be funded?

The answer is not simply "take the highest-scoring investments until the money runs out." That approach produces a list, not a portfolio. It ignores whether the money actually corresponds to available people. It ignores whether the resulting set is dangerously concentrated in one kind of bet. It ignores whether the mix has anything for tomorrow, or only things for today. Choosing the mix is the act of moving from a ranked set of values to a coherent, deliberate portfolio — one that the pool can actually afford, that the strategy needs, that the risk appetite permits, and that holds together as a set rather than a collection of individually-justified decisions.

This chapter is where the constraint meets the aspiration, and reality wins.

6.2 The funding line and what it means

Every portfolio has a pool. The pool has a bottom. The funding line is the point at which the money runs out, and everything above it is funded and everything below it is not.

In practice the funding line is not a clean horizontal cut through a ranked list. It is negotiated against several simultaneous constraints, applied in order. The starting point is the total available pool — the sum of money committed to portfolio investment for the period. This is not a number the portfolio sets; it is a number the portfolio receives. Working within it is the entirety of the job.

The first application of the pool is to existing commitments — investments already funded and running that have been re-entered into the comparison and are still above the line on forward value. They are not automatically continued; they earn continuation. But if they earn it, their remaining cost-to-complete is the first claim on the pool before any new investment can be considered.

What remains after covering the continuing investments — if anything remains — is available for new commitments. This remainder, measured against the demand for it, is where the real scarcity is felt. In a drifted portfolio of the kind described in Chapter 1, this remainder is often near zero, because the pool has already been committed to a set of incumbents that have been quietly auto-renewing for years. Making the remainder visible — and making explicit the choice between continuing an incumbent and starting something new — is one of the most valuable, and most uncomfortable, acts the demand and valuation stages enable.

The funding line is not a finding. It is a choice. Where it falls is a direct consequence of which incumbents were allowed back above it, and that choice deserves the same scrutiny as any new proposal.

6.3 Capacity versus demand

The funding line is drawn in money. But money is not the only constraint. People are, almost always, the binding one.

A portfolio funded to capacity on paper, but that exceeds the organisation's ability to staff and run it, does not run at the funded level. It runs at the staffed level, more slowly, with more collisions between teams, more half-started initiatives that cannot get the expertise they need, and more of the quality erosion that comes from people spread across too many things at once. The effect of exceeding the staffing constraint is not proportional expansion — it is the deterioration described in Chapter 1's dilution cost, only now deliberately caused.

"A portfolio is only as large as the people it can actually put to work. The rest is ambition parked in a spreadsheet, accruing interest in the form of delay and quality debt."

Capacity planning at the portfolio level is not the same as project-level resourcing. The portfolio does not schedule individuals or manage timesheets. It asks a coarser and more important question: do we have, in aggregate, the critical skills and capacities to run this set of investments at the funding level we are considering? The answer requires knowing which skills are genuinely scarce — the ones that appear across multiple investments and cannot easily be increased — and sizing the funded set against the realistic supply of those skills.

When the honest answer is that the funded set exceeds the staffed capacity, the portfolio has two options and only two: descope the set (stop or defer investments until capacity is available) or increase the capacity (hire, contract, or stop doing other things). Declaring the portfolio "aspirational" and proceeding anyway is not a third option. It is a decision to run a portfolio that will dilute its own investments, chronically miss its commitments, and produce less total value than a smaller, honestly-staffed set would have done.

6.4 Balance across four dimensions

Even a set of investments that fits within the money and the people is not yet a portfolio unless it is balanced. Balance is what separates a set of individually-justified investments

from a set that is coherent as a whole. There are four dimensions that require active management.

6.4.1 Risk and return

The selection criteria will have already filtered out investments whose individual risk is unacceptable. Balance adds a portfolio-level question: what is the overall risk profile of the funded set? A portfolio that is within the risk tolerance for each investment but where every investment is moderately uncertain at once is exposed to a bad year in a way that a genuinely balanced set — some safe, some speculative, some in between — is not. Conversely, a portfolio so heavily weighted to safe incremental work that there is nothing uncertain in it is guaranteed to under-invest in its own future.

The risk appetite set in Direction governs this: it specifies how much exposure to uncertainty the portfolio can carry, in aggregate, at any one time. Choosing the mix is where that appetite is actually applied — where the composition is inspected and, if necessary, rebalanced until the mix fits the stated appetite, not just the individual risk scores.

6.4.2 Time horizon

A portfolio invested entirely in near-term returns is safe today and fragile next year. A portfolio invested entirely in long-term bets is strategically interesting and operationally insolvent. Both are failures of balance.

Horizon balance asks: of the investments in the funded set, how much is creating value in the short term (within the next reporting period), how much in the medium term (one to three years), and how much in the long term (beyond three years, or uncertain)? The right proportions depend on the strategy and the organisation's current position, but the minimum requirement is that all three horizons have something — that the portfolio is not entirely consumed by the pressing at the expense of the important.

This is the second reason funding buckets earn their place, beyond the defensive one introduced in Direction. Buckets can explicitly protect each horizon by ensuring a minimum allocation to each — preventing the medium and long horizons from being crowded out in any single period by the always-compelling case for what needs to be done right now.

6.4.3 Strategic coverage

The strategy typically has several themes or objectives. A funded portfolio should advance each of them, or explicitly acknowledge which are deprioritised and why. A set that over-concentrates on one strategic objective and leaves others unaddressed is not a portfolio that reflects the strategy; it is a portfolio that reflects the strongest sponsor.

Strategic coverage asks: looking at the funded set as a whole, which strategic objectives are progressed, which are neglected, and is that neglect intentional? If it is intentional — if the strategy genuinely prioritises one objective above the others this period — it should be visible and stated. If it is accidental — if a strategic priority simply has no strong investment candidates and nobody noticed — that gap is information the portfolio should surface to the strategy, closing the upward link.

6.4.4 Run-the-business versus change-the-business

The most common balance failure is the slow erosion of investment in the future by the insatiable demands of the present. Run-the-business investments have the strongest and most immediate justification: operations will suffer if they are not funded, and operations are visible in a way that foregone growth is not. Change investments compete with each other for a share of what remains.

The funded set should have an explicit view of what proportion is running versus changing, and that proportion should be a deliberate choice, not an accident of which proposals were strongest. Where funding buckets are in use, this is managed through the bucket sizing. Where they are not, it requires a direct inspection of the funded set and a willingness to underfund a run-the-business investment — a genuinely uncomfortable act — if the run/change balance has tilted too far.

6.5 Concentration risk

Concentration risk is the fourth dimension given its own treatment, because it is the most commonly ignored and, when it lands, the most damaging.

A portfolio is concentrated when too many of its investments share a common dependency — a single technology platform, a key supplier, a single market, a specific regulatory assumption, a scarce skill set, or a strategic hypothesis that several investments are all betting on simultaneously. If that dependency fails — the platform doesn't scale, the supplier has a problem, the market assumption proves wrong — the portfolio does not lose one investment. It loses several at once.

Concentration risk is what turns a bad year into a crisis. Individual investment risk is managed by the selection criteria. Concentration risk is managed only at the portfolio level, by someone looking at the whole set and asking what they are all quietly depending on.

The check for concentration risk is a habit of mind, not a formula. Looking at the funded set, the question is: what would have to go wrong for more than one of these to fail simultaneously? The answer names the concentration. Whether the concentration is acceptable depends on the risk appetite and on the portfolio owner's view of how likely the shared dependency is to fail. High concentration on a stable, mature dependency is a different matter from high concentration on a new, unproven technology or a geopolitical assumption the strategy made a year ago and has not revisited.

Where concentration is found to be unacceptable, the portfolio has three responses: stop or reduce one or more of the concentrated investments, add an investment that diversifies the exposure, or explicitly accept the risk and note it for the next Review. The worst response — the most common one — is not to notice.

6.6 Funding buckets in action

With the balance dimensions in view, the role of funding buckets at the choosing stage becomes fully operational. Buckets, set in Direction, have defined the approximate proportions of the pool pre-committed to each category. Choosing the mix within a bucket means running the full comparison — forward value, capacity, balance — for the investments competing within that category, setting the funding line for that bucket, and stopping when either the money or the capacity runs out.

The sizing of the buckets is itself a choice — and a recurring one. Between periods, when the strategy is reviewed or the reference framework rebalanced, the bucket sizes are resized to match: a strategy that has shifted toward growth expands the change and innovation buckets; a strategy of stabilisation expands the run bucket. The resizing is a portfolio-level decision, not a per-investment one, and it happens before the within-bucket competition runs.

The protection buckets provide is most visible at the moment of competition: without them, every investment in the pool competes for the same money, and the safe, certain, near-term case reliably wins. With them, the innovation budget can only be taken by innovation candidates — a candidate that would lose against a run-the-business investment in open competition wins, within its bucket, against other innovation candidates that it is actually comparable to.

6.7 Making real trade-offs

The final act of choosing the mix is the one that distinguishes a portfolio that is managed from one that is merely administered: making the trade-offs explicit, visible, and owned.

A real trade-off is not "we cannot afford Investment A." It is "we are choosing Investment B instead of Investment A because, at this funding level with this strategy and these weights, B creates more forward value from this pool." The candidate below the funding line is not a failed proposal; it is the named cost of funding the one above it. That cost should be stated and visible, because it is the most honest account of what the portfolio is choosing not to do — and the portfolio's choices include its omissions.

"A portfolio decision that does not name what it is declining is a decision that has hidden half of itself. The funded set and the declined set are equally the portfolio's choices."

At every setting, from Lean to Enterprise, the test of a well-run choosing process is the same: can the people who made the decision explain, specifically, why each funded investment is in rather than the ones that are out — in terms of the framework, the weights, the capacity, and the balance? If the answer is yes, the portfolio can be

challenged and improved. If the answer is a general sense that these are the right things, the portfolio is being governed by intuition, which means it cannot be held to account.

6.8 How We Choose the Mix at three settings

Aspect	Lean	Managed	Enterprise
Comparison	Owner reviews the scored list against capacity and makes the call, visible on a single page	The board reviews a scored comparison table at the regular meeting	Formal comparison session with scored submissions, capacity confirmed, balance checked explicitly
Capacity check	Owner's knowledge of the team	A capacity-vs-demand view maintained each period	A formal capacity model by skill category, updated before each cycle
Balance check	Owner's judgement of the mix	A balance view across risk/horizon /rtb-vs-ctb reviewed at the board	Formal portfolio balance view, reviewed by the governing board before funding is confirmed
Concentration check	A deliberate question: what are these all depending on?	A dependency review at the comparison session	A formal concentration review as part of the cycle
Trade-off record	A note of what did not make the cut and why	A record in the demand log of declined candidates and the reason	A formal record, version-controlled, reviewed at each subsequent cycle to check whether circumstances have changed

With the mix chosen — funded above the line, staffed against real capacity, balanced across risk, horizon, strategic coverage and run-versus-change, and free of unacceptable concentration — the portfolio has its decisions. The next stage, Fund, is where those decisions become commitments: the money is released, the tranches are structured, and the rhythm of funding is established.

CHAPTER 7

Fund

Fund

7.1 From decision to commitment

Choosing the mix produced a funded set — a list of investments the portfolio has decided to back. The decision exists, but the money has not moved. Fund is the stage that makes the decision real: it releases the money, structures how it flows, and establishes the rhythm on which the portfolio will re-examine everything. A decision without funding is an intention. Funding is the commitment.

This stage is shorter than its neighbours, but it carries a structural weight that is easy to underestimate. The way money is released — the frequency, the conditionality, the staging — is the mechanism that governs how much flexibility the portfolio retains and how quickly it can change course when what it has funded turns out to behave differently from what was expected. An organisation that releases all its money at the start of the year, in a single tranche, to every investment in the funded set, has made a full-year commitment before the year's work has started. An organisation that releases money in short tranches, subject to review, retains the option to redirect it before the year is out.

The difference between these two is not a detail of treasury management. It is the difference between a portfolio that is re-decided on a rhythm and one that is decided annually and then simply executed.

7.2 Modes of funding

There are three basic approaches to releasing money, and most portfolios use a combination of them.

Rolling funding releases money continuously, in a regular cadence — monthly or quarterly — without a formal gate at each release. The investment is funded for the period; if nothing has changed, it is funded again for the next period without a specific review event. Rolling funding is fast, low-overhead, and appropriate for investments that are well-understood, reliably performing, and where the strategic context is stable. It is the natural funding mode for run-the-business investments: the organisation needs the lights on, and stopping to re-examine that choice every month is overhead without value.

Periodic funding releases money at the portfolio's regular re-decision cadence — typically quarterly — following a review of the whole funded set. Each period, the investments are re-examined, the pool is re-allocated against current priorities, and funding is committed for the next period only. This is the default mode for change-the-business investments: regular enough to provide working capital, infrequent enough to make each period's funding deliberate rather than automatic.

Gate funding releases money in stages, with a formal review and decision required before each stage can begin. The investment must demonstrate, at each gate, that it has earned

the next tranche — that the assumptions it was funded on still hold, that the value it is expected to create is still credible, and that continuing is better than stopping. Gate funding is the most resource-intensive approach and is appropriate for large, long-running investments where the risk is high enough, and the options wide enough, to justify the governance overhead.

The funding mode is a choice, and it should be chosen deliberately for each investment. Rolling for what is known and stable. Periodic for the active portfolio. Gated for what is large and uncertain enough to warrant formal scrutiny at each stage.

Most portfolios use all three modes simultaneously, with different investments on different rhythms. The portfolio cadence — the regular beat at which the overall funding picture is reviewed — governs the periodic mode and sets the outer limit for rolling; it does not require every investment to have a gate event at every cycle.

7.3 Cadence: the heartbeat of the portfolio

The portfolio's cadence is the regular rhythm on which funding is re-decided. It is one of the four dials, and it is the one with the most visible effect on how the portfolio feels to run.

A fast cadence — monthly, or even two-weekly — keeps the portfolio highly responsive and allows rapid reallocation when circumstances change. Its cost is overhead: each re-decision cycle requires some effort to run, and a frequent cadence can consume enough management attention to crowd out the actual work of managing the portfolio. Fast cadences suit Lean portfolios with a small set of investments and a high rate of change.

A slow cadence — annual, or longer — is low-overhead and stable, and it is the natural companion of the annual budget cycle in large organisations. Its cost is rigidity: a portfolio that can only re-allocate its money once a year is an organisation that is committed to the choices it made in October for the following twelve months, regardless of what the world does between January and September. In fast-moving environments, a twelve-month funding lock is not stability — it is brittleness disguised as process.

“Cadence is the portfolio's tempo. Too slow and it cannot respond. Too fast and it never stops responding long enough to actually fund anything properly.”

Between these poles, most portfolios find a quarterly cadence serves well for the bulk of their investments: frequent enough to catch significant changes in a reasonable time, infrequent enough to give investments a settled funding horizon within each period. The cadence is set as part of the operating model — the subject of Chapter 10 — and it is one

of the dials that can and should evolve as the organisation's portfolio management capability matures.

Within the cadence, there is an important discipline: funding stays steady between re-decision points. The portfolio commits for the period and does not disrupt running investments with mid-period changes unless the change is substantial enough to justify the disruption. This creates a working stability — the people doing the work can rely on the funding for the period — while preserving the portfolio's right to re-decide at the next cycle. Mid-period disruption is available but reserved for genuinely significant changes: a strategic shift, an investment whose performance has deteriorated materially, or a new opportunity whose value significantly exceeds what is currently running. Not for second-guessing this week what was decided last month.

7.4 Staged funding for uncertain bets: fund-to-learn

Uncertain investments — the ones flagged at the demand stage, valued as options in How We Value — require a fundamentally different funding structure. They cannot be funded in the same way as investments whose return is known and whose execution path is clear, because they are not that kind of investment.

The fund-to-learn structure releases money in tranches, each tranche funding one increment of learning, with a decision point at the end of each increment before the next tranche is released. The amount committed at each stage is the amount needed to answer the specific question that tranche is designed to answer — no more. The decision at each tranche review is not "is this on track?" but "what did we learn, and does what we learned make the next tranche worth funding?"

Fund-to-learn is not a phased delivery plan. It is a structured discovery process. Each tranche buys a piece of information; the decision to buy the next tranche is made on the basis of what the last one revealed.

The practical structure is simple. Each tranche has three elements: the learning objective (what specific question this tranche will answer), the amount (the minimum investment needed to answer it), and the decision criteria (what the answer needs to look like for the next tranche to be committed). A tranche that meets its learning objective — whether the answer is positive or negative — is a successful tranche. A tranche that fails to answer its question because it was not designed to has been mis-structured.

This structure has several important properties. First, it limits downside: the maximum loss on an uncertain bet is the sum of the tranches released before the point at which the learning made continuing obviously wrong. Second, it forces early learning: rather than designing the whole investment upfront and discovering near the end that a core assumption was wrong, the tranche structure is designed to test the most critical assumptions first, so the expensive later work only happens if the cheap early work said it should. Third, it creates legitimate stopping points: the end of each tranche is a governed moment at which not continuing is a fully acceptable and well-structured outcome — not a failure, but the correct response to a learning that said "not this, not now."

The fund-to-learn card — one of the templates in the annexes — captures this structure for each uncertain investment: the sequence of tranches, the learning objective and decision criteria for each, and the total envelope within which all the tranches sit.

7.5 The funding and commitment sheet

Every investment in the funded set, regardless of its funding mode, is documented on a funding and commitment sheet: the record of what has been committed, in what tranches, on what conditions, and with what decision rights.

The sheet is not a project plan. It does not track milestones or deliverables (that is delivery management). It records the portfolio's commitment: this investment receives this much money in this period, with this review date, and the portfolio's right to re-decide at that date. For gated investments, it records the tranche structure and the gate criteria. For fund-to-learn investments, it records each learning tranche and its decision criteria.

The aggregate of the commitment sheets is the portfolio's current allocation of the pool. The total committed must not exceed the total pool. The review dates, taken together, constitute the portfolio's cadence — the heartbeat described above.

7.6 Fund at three settings

Aspect	Lean	Managed	Enterprise
Funding mode	Rolling for known work; owner's judgement on uncertain bets	Periodic funding at the quarterly cadence; gated for large or uncertain investments	Rolling for run, periodic for change, gated for transformation; fund-to-learn governed separately
Cadence	As often as the owner needs; can be continuous	Quarterly, with a formal review meeting	Formal cycle aligned to governance calendar; may be quarterly or semi-annual
Commitment sheet	A shared list of what is committed and for how long	A short structured sheet for each investment	A formal funding record, version-controlled, ratified by the governing board
Fund-to-learn	Owner sets the tranche size and the learning question	A standard template, reviewed at the intake gate	A governed pathway with formal tranche approvals and a learning review between tranches
Mid-period change	When the owner judges it necessary	When the change is material enough to bring to the board outside the normal cycle	A formal exception process with defined thresholds for what constitutes a material change

At the Lean end, the risk is treating all funding as continuous and unconditional — releasing money to everything in the set and never structuring the re-decision moment. The portfolio then runs on the implicit assumption that the funded set is right until something obviously goes wrong, which is the same pattern as the drifted portfolio but now with a process label on it. The discipline at Lean is to keep the review moment real: to genuinely ask, at the cadence, whether each investment still earns its place.

At the Enterprise end, the risk is the opposite: so many gate events, so many approval layers, and so much documentation that the portfolio's ability to commit quickly has been traded away for a governance architecture that moves money in geological time. The test of a well-designed funding process at any setting is not whether it has sufficient controls. It is whether it can fund the right work at the speed the organisation needs.

With funding committed — the money released, the tranches structured, the cadence established — the portfolio's chosen mix is now running. That is when the second job of the portfolio begins in earnest: making it work together. That is Steer.

CHAPTER 8

Steer

Steer

8.1 The second job begins

Funding committed the money. The investments are running. And this is precisely the moment when the gap between the portfolio as chosen and the portfolio as it actually behaves begins to open.

The gap is not a surprise. It is a structural feature of any complex set of investments running simultaneously through a shared organisation. The selection criteria assessed each investment on its own merits – its value, its strategic fit, its individual risk, its cost. They did not, and cannot, assess the investment as a member of a set: how it will interact with the others, what it will need from shared resources, which of its assumptions depend on something another investment is also depending on, what it is quietly blocking. The choosing was about which investments. Steer is about what happens when all of them are trying to run at once.

This is where the portfolio's second job – making the chosen mix work together – becomes concrete. And it is where the method's claim to do something useful, beyond the obvious discipline of choosing carefully, stands or falls.

8.2 What Steer is for

Steer is not a monitoring process. Monitoring – collecting status from running investments, summarising it upward, presenting a colour-coded dashboard to a governance forum – is something most portfolios already do, at considerable effort and with limited return. Monitoring asks: how are things going? The answer it receives is the answer the investment wants to give, which is the answer most likely to leave its funding undisturbed. The eighteen-month green status that turns red the day the investment is cancelled was not produced by a corrupt monitoring process; it was produced by a monitoring process that had no mechanism other than self-reporting.

Steer is not a monitoring process. It is a decision-forcing process. Its purpose is to surface the problems that no single investment owner can or will surface alone, and to produce decisions – not reports.

The distinction is not subtle. A monitoring process produces information. A Steer process produces choices: this team waits while that one gets the scarce resource; this dependency is unblocked by changing the sequence; this investment is paused because what it is blocking is more valuable; this duplication is eliminated by merging two workstreams. Information is an input to Steer. Decisions are its output. An investment that presents its problems in a Steer forum and receives only acknowledgement has not

been steered. A problem acknowledged is not a problem resolved.

8.3 The anatomy of a Steer problem

The problems Steer exists to solve fall into a small set of categories, each with a distinct character and a distinct kind of resolution.

8.3.1 Dependencies

A dependency exists when one investment's ability to proceed depends on an output, decision, or action from another investment or another part of the organisation. Dependencies are the most common source of delay in a running portfolio, and they are invisible to the individual investment owner: the owner of Investment A can see that A is waiting for something; they cannot see, from inside Investment A, why Investment B has not yet produced it, whether B has been reprioritised, or whether the decision that would unblock both of them is sitting in someone's inbox waiting for a meeting that has not been scheduled.

Managing dependencies at the portfolio level means maintaining a live picture of what is waiting for what, across the whole set, and treating each significant dependency as a portfolio problem rather than a bilateral negotiation between the two teams involved. The bilateral negotiation is the default; it is slow, it escalates only when the teams have already exhausted their own options, and it produces a settlement based on the relative seniority of the sponsors rather than the relative value of the investments. Portfolio-level dependency management asks which resolution serves the whole mix best, and makes that the decision.

8.3.2 Scarce resource contention

Some resources are scarce across the portfolio: a specific technical capability, a subject-matter expert, a test environment, a regulatory specialist, a technology platform that can only support one major change at a time. When multiple investments need the same scarce resource simultaneously, the resource must go somewhere – and where it goes is a portfolio decision, not an individual investment's decision.

The common failure is to leave this decision to the people who can see only their own piece of it: the resource is pulled in multiple directions, gives a little to each claimant, and contributes fully to none. The result is the dilution described in Chapter 1, now operating within the running portfolio rather than across the funded set. Steer resolves resource contention by making the competition visible – the resource-contention view shows which investments are competing for the same resources – and by making a clear, portfolio-level decision about the allocation. One investment gets the resource. The other waits, or is rescheduled, or is partially paused. The decision is uncomfortable and specific and necessary.

“Shared-resource contention resolved by the portfolio is uncomfortable for someone. Shared-resource contention not resolved by the portfolio is uncomfortable for everyone, indefinitely.”

8.3.3 Sequencing

Not every investment can run at full pace simultaneously, even if each is independently funded and staffed. Some create dependencies for others; some compete for the same window of organisational attention; some are building on foundations that another investment is laying. The optimal sequence for the running portfolio — which investments run at full pace in which period, which slow down to allow others to complete a prerequisite, which defer their peak resource use until the constraint has cleared — is not the same as the schedule each investment would choose for itself.

Sequencing decisions belong to the portfolio, not to the individual investment. An investment asked to slow down so that another can proceed faster is not being managed badly; it is being managed as part of a set. The Steer process owns the sequencing view — the picture of how the investments' timelines interact — and makes the adjustments needed to keep the whole set moving at the pace the pool and its capacity can support.

8.3.4 Duplication

Two investments building the same thing is waste. Three investments each building a slightly different version of the same capability, in different parts of the organisation, because no one looked across the whole set, is the kind of waste that compounds: the three versions are incompatible, the organisation ends up running all three in production, and the integration problem arrives later as a crisis.

Duplication is almost never visible from inside either of the duplicating investments. Each owner has a legitimate need; each is building a legitimate solution to that need; neither can see that the other exists, or that the capabilities they are building will turn out to overlap. The Steer process finds duplication by looking across the whole set with a single view — the portfolio dashboard and the dependency map together show where workstreams are converging on similar territory — and resolves it before it hardens into a structural problem.

Resolving duplication rarely makes anyone happy in the short run. It means telling one investment — or both — to change what it is doing, to use what the other is building, to merge, or to stop. These are the conversations that no individual sponsor can initiate without appearing to attack a colleague's programme. The portfolio forum is the right place to have them, precisely because it holds the mandate to look at the whole.

8.3.5 Blockers the investment cannot clear alone

Some of the problems a running investment encounters cannot be resolved by the investment itself. A decision that requires a stakeholder outside the investment's sponsor chain; a dependency on a supplier relationship owned elsewhere in the organisation; a

regulatory query that requires a response from the legal function that the investment's programme manager cannot compel. These blockers are not failures of the investment — they are failures of the interfaces between the investment and its environment, and those interfaces are the portfolio's to manage.

The decision and escalation log is the mechanism: a running record of the blockers that have been identified, who is responsible for resolving each one, and by when. Items on the log are not problems awaiting resolution; they are commitments to specific resolutions, owned by specific people, with specific dates. A blocker that has been on the log for three consecutive Steer cycles without progress has not been escalated — it has been catalogued. Cataloguing is not Steer.

8.4 Steer forces decisions

The thread running through all four problem categories is the same: each requires a decision that is uncomfortable in some dimension — who waits, who gets the resource, whose timeline changes, whose workstream is modified or stopped — and each decision is one that no individual investment owner can make alone, because the decision's consequences fall on the whole set. This is the portfolio's second job in its most operationally demanding form.

A Steer process that does not produce these decisions is not a Steer process. It is a status collection exercise — valuable, perhaps, as information, but not as portfolio management. The test of a well-run Steer process is not the quality of its dashboards. It is whether, when two investments need the same architect, the portfolio says clearly that Investment A gets the architect this quarter and Investment B waits — because A is on the critical path that unlocks B's market — and whether that decision is implemented, not merely noted.

The measure of a Steer process is not the quality of its status pack. It is the number of decisions it produces, the speed with which those decisions are implemented, and whether the consequences are absorbed by the right investments rather than deferred until they become a crisis.

Producing decisions requires decision rights — clarity about who has the authority to make which calls — and decision rights require the operating model. Chapter 10 develops the full decision-rights picture. What matters here is the principle: Steer must have the authority to make the decisions it needs to make, or it degrades into a forum where problems are named and nobody acts.

8.5 Feeding hard truths to Review

Steer is where the optimism of selection first meets the friction of reality. Every investment was chosen on a view of its forward value and its delivery risk that was formed in advance of the work. As the work proceeds, that view is updated — sometimes

positively, more often not — and the updated view belongs in the portfolio's next re-decision, which is Review's job.

Steer's contribution to Review is the honest account of what is actually happening: which investments are progressing as expected, which are struggling, which have surfaced information that changes the assessment of their forward value, and which have blocked or been blocked in ways that have real cost implications. This is the hard-truths function, and it is structurally important that Steer, not the investment itself, is the primary source of it.

An investment asked to self-report to a Review forum has a structural incentive to report charitably. It has a sponsor whose commitment is attached to the investment's success, a team whose work is being assessed, and a set of assumptions it was funded on that it has every reason to protect. Steer's view is not self-reported; it is the portfolio's view of the investment from outside, synthesised from the dependency map, the resource-contention picture, and the escalation log. When these two views diverge — when the investment's self-report says "progressing well" and the Steer process says "blocked on three dependencies, consuming a scarce resource needed elsewhere, and six weeks behind the sequence it was funded on" — the portfolio has the information it needs, and Review can act on it.

8.6 The Steer toolkit

Four instruments support the Steer process. They do not replace judgement; they provide the visibility without which the judgements cannot be made.

The portfolio dashboard and health view is the at-a-glance picture of the running set: the investments, their status against the funding commitment, any flagged issues, and the overall health of the mix. It is the starting point for every Steer meeting. Its purpose is not to report that everything is fine — if everything is fine, the meeting can be short. Its purpose is to surface quickly the investments that need the Steer process's attention.

The dependency map shows what is waiting for what, across the whole running set, with the critical-path implications of each dependency. It answers the question: if this dependency is not resolved by this date, which investments are affected, and how? Maintained as a live document rather than a one-time snapshot, it is the instrument that makes sequencing decisions discussable.

The decision and escalation log records every blocker, decision, and escalation that has emerged from the running portfolio: what the issue is, who is responsible for resolving it, by what date, and what happened. It is the accountability record of the Steer process. An issue on the log is a commitment, not an observation.

The resource-contention view shows where the portfolio's scarce resources are currently allocated across the running set, and where contested demands for them exist. It makes the invisible visible: without it, each team knows only its own resource situation; with it, the portfolio can see the collision and decide who gets priority.

8.7 Steer at three settings

Aspect	Lean	Managed	Enterprise
Steer forum	Regular check-in with the owner and key leads; can be brief and informal	A standing Steer meeting on the cadence, with representation from each running investment	A formal Steer forum with defined attendees and decision rights; may have sub-forums by portfolio tier
Dashboard	A shared list or board visible to the owner	A structured portfolio health view, updated before each Steer meeting	A governed dashboard, updated continuously, with defined health metrics and thresholds
Dependency map	The owner's working knowledge, made visible when needed	A maintained dependency map, reviewed at each Steer meeting	A formal dependency register with impact analysis and critical-path tracking
Resource contention	Owner resolves directly, visible to all involved	A resource-contention view reviewed at the Steer meeting; decisions made in the meeting	A formal resource-allocation view; significant contention escalated to the governing board
Escalation log	A shared list of blockers and who is resolving them	A structured log reviewed at each Steer meeting, with owners and dates	A formal escalation register with SLA tracking for resolution

At the Lean end, the failure mode is running Steer as a social update rather than a decision forum. When the portfolio is small enough that everyone knows everyone and the relationships are good, problems surface informally — which is fine — but the decisions are also made informally, which means they are sometimes not made at all, or made inconsistently, or made without the portfolio picture that would produce a better outcome. The discipline at Lean is to keep the decision moment real, even if the surrounding process is minimal: the resource goes somewhere specific, the dependency is unblocked by a specific action, the blocker is owned by a specific person with a specific date.

At the Enterprise end, the failure mode is the inverse: a Steer process so procedurally elaborate that it consumes more management attention than the problems it is solving, produces reports that take longer to prepare than to read, and escalates at a speed that makes its resolutions irrelevant by the time they arrive. The portfolio dashboard that takes three days to produce weekly is not a decision tool. Enterprise Steer needs the substance of every category above — visible dependencies, resolved contention, cleared blockers, sequenced work — at a tempo the organisation can actually respond to.

Steer does not end at a pre-determined point. It runs continuously while the portfolio runs, updating its views and producing its decisions as each cycle of the cadence arrives. What it is preparing for — at each cadence — is Review: the moment at which the whole running set is re-opened, re-valued, and re-decided. That is the next chapter.

CHAPTER 9

Continuous Re-prioritisation and Reallocation

Continuous Re-prioritisation and Reallocation

9.1 The loop comes back around

Steer has been running. The investments have been moving — some well, some not, some revealing information about themselves that nobody had in advance. The dependency map has entries. The escalation log has owners. The portfolio dashboard shows the picture. And now the cadence arrives at the moment the whole method is built toward: the periodic re-opening of the question of what should be funded.

Review is the stage where the portfolio re-decides. Not every investment, every cycle, in a full formal process — that would be an organisation that does nothing but decide about deciding. But the whole live set is re-examined, against all available information, on the rhythm the cadence sets, and the money moves to follow the value wherever it now is. This is what makes the portfolio a living thing rather than a list that was approved and defended: the re-decision is structural, periodic, and built in. It is not a sign of failure when the mix changes; it is the system working.

Review is not an audit. It is not a performance tribunal. It is the portfolio asking, again, the same question it asked when it chose the mix: given everything we now know, and the limited money we have, is this still the best set of things to be funding?

This chapter is one of the two heaviest in the book, as flagged in the introduction, because it carries the most concentrated weight of the method. It is where the principle of forward value is tested against real investments that have real histories. It is where the politics of stopping work are most acute. It is where the learning that the portfolio has gathered feeds back into the strategy that created it. And it is where the promise of continuous re-allocation — the distinctive claim of the method — is either kept or quietly abandoned.

9.2 Re-judging the live set

The comparison that Review stages is not a comparison between the live investments and some abstract standard. It is a comparison between the live investments and each other, and between the live investments and the new demand waiting to be funded. Everything is on one table.

This is the single discipline that most distinguishes a managed portfolio from a drifted one. In the drifted portfolio, the incumbents are not on the table. They are the table: the fixed context within which new arrivals are considered, approved in small increments, or declined because there is no room. The managed portfolio puts the incumbents on the

table alongside every new candidate and asks the same questions of all of them: what is the forward value, against what it will still cost? How does that compare to the alternatives? Is this the best claim on the money from today?

“The incumbent's advantage is not momentum or history. It is only the forward value it still offers. Once that is no longer the best available claim on the money, the incumbent's time is done.”

Mechanically, the re-judgment works the same way as the initial choice: each running investment updates its One-Pager to reflect current knowledge, the valuation models from How We Value are applied to the updated numbers, and the resulting forward values are placed alongside the new candidates from the demand pool. The reference framework and the strategy-set weights are applied consistently across both. The funding line is then drawn at the point where the money runs out, and it falls where it falls — regardless of whether the investment above it is new or old, well-established or recently started.

The practical implication is that running investments must maintain their case, not just their momentum. The team that has been executing well but has not updated its forward value picture in six months is in a weaker position at Review than the team whose forward value estimate is current, specific, and honestly assessed. Performance is visible in the numbers; the numbers go into the comparison.

9.3 Moving money on forward value

The re-judgment produces a ranking. Investments above the line continue; investments below the line face a decision. But that decision is not the simple cut that a ranked list implies. Moving money from one investment to another is a real action with real costs — transition costs, team disruption, contract consequences, reputational effects — and those costs are part of the calculation.

The governing principle is Principle 6: re-decide on a rhythm, not on impulse; fund stays steady within a period; you re-judge everything between periods; you only move money when the difference is big enough to be worth the disruption. A marginal difference in forward value between two investments does not justify a reallocation; the transaction costs of moving the money may exceed the value gained. A significant difference — one that is both clear and durable — does justify it, and in that case the money should move, regardless of how uncomfortable the movement is.

Reallocation is not disruption for its own sake. It is the mechanism by which the portfolio keeps its promise that the money follows the value. Not reallocating when the value has clearly moved is a decision to waste the difference.

This is the mechanism that reverses drift. Drift accumulates when money continues to flow to work that no longer earns it, because no one has made the comparison and named the gap. Review makes the comparison structural and the gap visible. Reallocation is what the visibility requires.

9.4 Tracking benefits after delivery

The review of running investments addresses the work that is not yet done. But the portfolio's accountability does not end when an investment crosses the delivery line. A piece of work that was funded on a claim of value must be held to that claim — and the value almost always arrives, if it arrives at all, some time after the delivery is complete.

This is the benefits-after-delivery discipline, and it is one of the most widely neglected in portfolio practice. The common pattern is: investment delivers at go-live; the portfolio records the delivery; the investment is formally closed; the benefit claim is filed and not revisited. Eighteen months later, when the benefit was supposed to have arrived, the portfolio is funding twelve new things and nobody has checked whether the previous investment's benefit actually materialised.

The neglect is not accidental. It is structurally produced by the incentives: the sponsor who championed the investment has moved on, the team has been reassigned, and the budget holder who would have to account for the missed benefit is not the same person who made the promise. Benefits review after delivery requires someone to own the accountability beyond the delivery moment, and most governance structures do not create that owner.

The method's answer is the value and benefits tracker: a live record that follows each investment from the point at which it claims a benefit through to the point at which that benefit should have materialised, with a check at that point. The check asks: did the value arrive? If not, why not, and what does that tell us about the assumptions that justified the investment? The answers feed directly into Direction — they are evidence about what the portfolio's criteria and weights are actually selecting, and whether those selections are performing as expected.

“Measuring benefits only at go-live is measuring intention, not value. The value is what happens in the organisation after the delivery, not on the day of it.”

The benefits tracker is uncomfortable to run, for exactly the same reasons that stopping work is uncomfortable: it attributes accountability, sometimes for failures, to people who would prefer the question not be asked. That discomfort is the point. A portfolio that never checks whether its past choices actually delivered what they promised is a portfolio that cannot learn — and one that will keep funding the kinds of investments that have been quietly failing to deliver, because no one ever looked.

9.5 Stopping work

The most important — and least comfortable — act in the Review stage is stopping an investment that no longer earns its place.

The decision to stop is not a failure. It is the method working correctly. An investment that was funded on assumptions that have since proved wrong, or that offered forward value that now looks less compelling than the alternatives, should stop. Stopping it is not a punishment for the team that ran it or the sponsor who championed it; it is the portfolio correctly identifying that the money would serve the organisation better elsewhere, and acting on that finding. The fact that it is understood this way in theory and almost never experienced this way in practice is one of the central challenges of portfolio management.

The discipline of stopping begins before the investment is funded. Principle 12 requires that the rules and stop criteria are agreed in advance, before any investment is on the table and before any sponsor is defending it. A stop criterion is a specific, pre-agreed condition — a test of the assumptions the investment was funded on — whose failure triggers a formal re-decision. Not an automatic stop; a decision. But a decision that is not fighting against the assumption that continuing is the default.

Stop criteria agreed before the investment starts are honoured when the facts arrive. Stop criteria proposed when the investment is already in trouble are fought by everyone whose career, bonus, or reputation is attached to the investment's continuation.

Without pre-agreed stop criteria, the stop decision must overcome not only the investment's actual merits but the entire political apparatus that any significant investment accumulates: the sponsor's sunk emotional capital, the team's legitimate attachment to their work, the stakeholders who have built plans around the delivery, and the governance forum that approved the investment and is reluctant to admit that the approval was wrong. These forces are not illegitimate; they are human. But they are forces that push in the same direction — continuation — regardless of the forward value, and without pre-agreed criteria to counterbalance them, they reliably win.

9.6 The politics and the biases

The forces that resist stopping work are worth understanding specifically, because understanding them is the first step to designing a process that can survive them.

Sunk cost bias is the most pervasive. The money already spent on an investment is irrelevant to its future funding, as established in the canon. But it does not feel irrelevant. It feels like a reason to continue — as though stopping now would mean the spent money was "wasted," while continuing might yet redeem it. This is an error. The money is already spent whether the investment continues or stops; the only relevant question is whether continuing or stopping produces more value from today. Sunk cost bias turns a forward-looking decision into a backward-looking one, and it does so silently, in the framing of the conversation, before any numbers are on the table.

Escalation of commitment — sometimes called the "Concorde fallacy" — is the tendency to invest more precisely because so much has already been invested. Each additional tranche of money is justified by the previous one, producing a chain in which the investment continues long past the point at which it would have been declined if judged fresh. The method's defence is forward value: if the investment would not be funded now, on its current forward value, against the alternatives now in the pool, it should not continue being funded just because it has been.

Optimism bias is the investment's natural tendency to present its own prospects charitably. The team working on a difficult programme believes, in most cases, genuinely and not dishonestly, that the next quarter will turn the corner, that the technical problem will be solved, that the market will respond. This belief is often unfounded, not because the team is dishonest, but because proximity to the work makes its obstacles look temporary and its successes look predictive. The portfolio's external view, provided by Steer, is the partial corrective; the pre-agreed stop criteria are the structural one.

Sponsor protection is the most political of the forces. A senior sponsor attached to an investment is not a neutral assessor of its forward value. Their professional identity, their relationships, and sometimes their compensation are bound up in its success. A portfolio process that makes the stop decision dependent on the sponsor's agreement is a portfolio process that will very rarely stop anything its most powerful sponsors have backed. The process must make the decision on the forward value, and it must have the authority — and the governance structure — to make it stick.

"The investment that needs another ten percent — and has needed exactly another ten percent for three consecutive years — is not almost finished. It is a sponsor protection story told in small increments, each plausible on its own, adding up to a decision that was never made."

9.7 Designing for survivable decisions

The response to these forces is not to be tougher, or to demand that sponsors be more objective, or to exhort the portfolio forum to "focus on the data." These responses are well-intentioned and ineffective. The response is structural: design the process so that good decisions can survive the politics, rather than relying on the politics to produce good decisions.

The structural elements are: pre-agreed criteria, set before the investment starts and encoded in the funding and commitment sheet; a Review forum with the authority and the mandate to act on the criteria when they are met; a reallocation process that is visible and explained, so that the decision is understood rather than resented; and a communication discipline that names the facts — forward value, updated costs,

alternatives in the pool — rather than avoiding them in the hope that the conversation will be gentler.

None of these elements makes the stop decision easy. They make it possible: possible to reach, possible to explain, and possible to implement without being reversed by the next governance meeting at which the sponsor appears.

9.8 Benefits after delivery and the feedback to strategy

When an investment stops — whether at completion or by the portfolio's decision — its claims become testable. What it promised, in the business case and in the forward value estimate, either arrived or did not. The tracking of those outcomes is not a retrospective exercise in accountability for its own sake; it is the most direct input the portfolio has into the strategy that funded it.

If a class of investments that was systematically rated highly — strong strategic fit, good forward value, acceptable risk — has consistently failed to deliver its claimed benefits, that is information about the reference framework's validity. The criteria, or the weights, or the valuation model that produced the high ratings, has been selecting for something that does not perform. Direction needs to know this. The upward link between the portfolio and strategy — established in the canon, exercised through Review — is the channel through which this learning travels.

This is the most concrete expression of the two-way strategy link. Strategy does not merely send instructions downward into the portfolio. The portfolio sends evidence upward: evidence about what kinds of investment actually create value, what assumptions turn out to be wrong, what risks materialise and what risks prove illusory. An organisation that runs this feedback loop honestly — that allows Review's findings to change the weights and sometimes the criteria — has a portfolio that becomes progressively better at choosing. An organisation that treats the portfolio as a pure executor of strategy, sending results upward only as performance data and never as learning, will repeat the same selection errors indefinitely.

The upward feedback from Review to strategy is not a report on how well the portfolio executed the plan. It is evidence about whether the plan was right, and about what to change in the framework that chose the investments that tested it.

9.9 Continuous re-prioritisation at three settings

The re-judgment of the live set happens on the cadence, at every setting. What differs is the formality, the frequency, and the depth of the comparison.

Aspect	Lean	Managed	Enterprise
Re-judgment trigger	On the owner's cadence; continuously if needed	At the regular portfolio board meeting each period	At the formal review cycle, with a full submission from each investment
Comparison table	A current list of all live and incoming investments with forward values	A structured comparison on the standard scoring template	A formal comparison pack, with updated valuations and Steer's input, reviewed by the board
Benefits tracking	Owner tracks whether value arrived; checks at agreed milestones	A value and benefits tracker maintained for each investment, reviewed at each cycle	A formal benefits register, version-controlled, with post-delivery review at agreed intervals
Stop process	Owner makes the call, records the reason	A stop-or-continue decision at the board, against the pre-agreed criteria	A formal stop decision, ratified by the governing board, with a communication plan and reallocation record
Reallocation	Money moved immediately on the owner's decision	Moved at the next funding release, with the board's ratification	Moved at the next formal funding cycle, via the gated release process
Feedback to strategy	Owner notes what worked and what did not; informs the next Direction conversation	A lessons log reviewed at each cycle; significant findings escalated to the strategy team	A formal learning review feeding the next strategy refresh and the next Direction re-set

At the Lean end, the risk is that re-judgment becomes rare in practice — not because the owner intends to ignore it, but because the steady pressure of running work crowds out the periodic step back that would reveal which of it should stop. The discipline at Lean is the cadence: a commitment to the regular moment of comparison, even when the running work seems to be going well enough, because "going well enough" is not the same as "best use of the money."

At the Enterprise end, the risk is the reverse of what might be expected: not too little re-judgment, but re-judgment that is too formal and too slow to catch the investments that are quietly eroding. A six-month formal review cycle means that an investment can deteriorate for three months before the review sees it, and then wait three more for the next cycle to do anything about it. Enterprise portfolios must supplement the formal cycle with interim signals — Steer's hard truths, the escalation log, the benefits tracker — so that significant changes in forward value can be acted on between formal review points.

Review is the end of one turn of the loop. It is also the beginning of the next. The learning it gathers, the reallocation it executes, and the signals it sends back to Direction are the inputs that make the next round of choosing better than the last. The loop does not end here; it loops. That is the point.

PART

Part III — Running It For Real

CHAPTER 10

The Operating Model

The Operating Model

10.1 The method needs a home

The Investment Loop is a model — a description of what needs to happen, in what order, on what rhythm. A model does not run itself. It needs people with the right authorities, a structure that supports their decisions, and a set of explicit choices about how formally and how often the loop turns. The operating model is that home: the infrastructure of roles, rights, and settings within which the method lives in a real organisation.

This chapter is not a re-teaching of the stages. The stages have been covered. This is the question of who does them, how they are governed, what the decision rights are, and how an organisation calibrates the method to its own scale and maturity. It is also the question of how the method scales beyond a single portfolio — into the nested, tiered structures of a large enterprise — and how the settings evolve as capability grows.

10.2 Roles

The method requires a small number of roles, each with a clear and non-overlapping purpose. The precise titles vary by organisation; the functions do not.

The portfolio owner is the person accountable for the portfolio's choices: what is funded, what is not, how the mix is managed, and what is stopped. At the Lean setting this is a single person — a founder, a product director, a head of a division — whose accountability is personal and complete. At larger settings, accountability may be shared, but it must be locatable: there must be a specific person, or a governing body with a named chair, who carries the responsibility and cannot successfully diffuse it.

The portfolio board (at Managed and Enterprise settings) is the governing body that makes the funding decisions: the Direction setting, the comparison and choice, the Review. It is not an advisory group. It is a decision-making body with the authority to commit and to stop. Its composition should reflect the domains and stakes it governs, and its authority should be unambiguous — a board that lacks the authority to stop a significant investment is not a portfolio board. It is a presentation venue.

The portfolio office (at Managed and Enterprise settings, sometimes at Lean) is the function that runs the mechanics of the loop: managing the demand log, preparing the comparison, coordinating the Steer process, maintaining the benefits tracker, and keeping the cadence. The portfolio office does not make funding decisions; it enables and services them. The distinction between owning the decision and enabling it is important: a portfolio office that has drifted into making the funding choices, or that has accumulated enough process authority to effectively determine the outcome before the board meets, has broken the governance model. The decision must be made by the right people, with full visibility of the options.

The investment sponsor is the senior individual accountable for a specific investment: for making the case for it at intake, for the accuracy of the One-Pager and the forward value estimate, and for the investment's performance against what it promised. The sponsor is not the portfolio's customer; they are one of its principals. A sponsor who is not held to the accuracy of their forward value claims, or who can propose investments without being accountable for their outcomes, is a sponsor who will consistently over-promise.

The investment lead (or programme manager, or product manager) is responsible for the delivery and management of a specific investment — the inside of it, which is delivery management and outside the portfolio's scope. The investment lead reports to the sponsor, not to the portfolio, on everything internal to the investment; they report to the Steer process on the interfaces between their investment and the rest of the portfolio.

The portfolio's governance is about the relationships and decisions between investments. The sponsor and investment lead own the inside of the investment. Conflating these produces a portfolio function that is trying to run the work — which is the boundary violation the method prohibits.

10.3 Decision rights

Knowing who plays which role is necessary but not sufficient. What matters in practice is knowing who decides what — and at what threshold a decision must be escalated to a higher authority. Decision rights are the explicit statement of that mapping.

A well-designed decision-rights framework answers three questions for every significant decision the portfolio makes. Who decides? Is this the portfolio owner's call, the portfolio board's, or a sub-portfolio board's? At what level? What size, complexity, or strategic significance of investment triggers a requirement for board ratification rather than owner discretion? On what basis? What information and process must the decision be based on — the reference framework, the comparison table, the stop criteria — to be valid?

Without explicit decision rights, authority is ambiguous, which means it is contested. Sponsors whose investments are declined appeal to the next level. Portfolio boards approve investments they have not really compared. Owners stop investments without the board's awareness and produce a governance crisis. The decision-rights framework is not a bureaucratic indulgence; it is the architecture that lets the loop run without producing a political conflict at every decision point.

Decision	Lean	Managed	Enterprise
New investment approval	Portfolio owner	Portfolio board	Portfolio board, or sub-board within agreed thresholds
Continuation / stop	Portfolio owner	Portfolio board	Portfolio board; significant stops ratified by executive sponsor

Decision	Lean	Managed	Enterprise
Reallocation between periods	Portfolio owner at any point	Portfolio board, triggered by material change	Formal exception process; executive escalation above defined thresholds
Bucket resizing	Portfolio owner at each period	Portfolio board at each period	Strategy function with board ratification
Reference framework re-weighting	Portfolio owner	Portfolio board, with strategy input	Strategy function; board ratification; version-controlled

10.4 Governing the prioritisation process itself

The portfolio board governs the investments. Who governs the portfolio board? This question, often unasked, matters enormously. A board that operates without any process governance — no standards for what it reviews, no requirements for the quality of the comparisons it makes, no accountability for the consistency of its own decisions — is a board that will drift from the method toward intuition and then toward politics, without anyone formally choosing to make that move.

Governing the process means: the rules for what constitutes a valid submission to the demand stage; the standards the comparison must meet before the board decides; the cadence the board is committed to, and the consequences of missing it; the criteria for what triggers a between-cycle escalation; and the mechanism for reviewing whether the portfolio's decisions are producing the outcomes they claimed.

At Enterprise scale, this is often a formal governance charter, ratified by the executive. At Lean, it is the owner's own standards, made explicit. At Managed, it sits somewhere between: a set of working agreements the board has made about how it operates, revisited occasionally, and available to be invoked when someone is behaving inconsistently with them.

10.5 Setting the dials

Chapter 2 introduced the four dials — cadence, who decides, formality, funding rigidity — and the three named settings they resolve into. The operating-model chapter is where those dials are actually set.

The Dial-Setting Worksheet (Annex G) guides this process. It is not a self-assessment against best practice; it is a deliberate choice about which setting genuinely serves the organisation's current situation — its size, its portfolio complexity, its governance maturity, and the speed at which its environment changes. The choices are interdependent: a portfolio at Lean cadence (fast) cannot realistically run Enterprise-level formality (slow), not because the combination is prohibited, but because it will break under the contradiction in practice.

The starting point is the organisation's current reality, not its aspiration. The question is not "what setting would make us look most sophisticated?" but "what setting can we actually run well, consistently, and honestly?" A portfolio that is operating at Managed on paper and Lean in practice — because the formal processes are gamed, bypassed, or ignored — has chosen its setting incorrectly. The method works at Lean just as rigorously as at Enterprise; the rigour is in the discipline of the questions, not the weight of the documentation.

The dials evolve. A portfolio that starts at Lean and grows — more investments, more money, more stakeholders, higher stakes — naturally requires more governance, and the dials should move to Managed and eventually Enterprise in response. That evolution should be deliberate and acknowledged, not unconscious and contested. The portfolio office (when it exists) typically triggers the conversation about dial adjustment, because it observes the friction points — decisions that cannot be made at the current formality level, conflicts that the current authority structure cannot resolve — that signal a setting change is needed.

10.6 Portfolios of portfolios

Large organisations are not single portfolios. They are portfolios of portfolios: a corporate portfolio that allocates across divisions, divisional portfolios that allocate across functions, functional portfolios that allocate across teams. The nesting can be two layers or five, depending on scale.

The method scales into this structure directly: each portfolio in the hierarchy is running the same Investment Loop, with the same stages, on its own cadence and at its own setting. The higher-level portfolio treats the lower-level portfolio's total funding commitment as an investment in its own pool — it allocates to the division, not to the individual investments within the division. The division's portfolio allocates within its envelope.

At each tier, the portfolio allocates to the level below, not into it. The corporate portfolio funds the divisional portfolio, not the divisional programmes. Each tier has full authority within its envelope and no authority beyond it.

The governance connections between tiers are principally three. First, Direction: the corporate portfolio's reference framework and strategy-set weights flow down as constraints on the divisional portfolio's own Direction. The divisional portfolio may have local criteria and weights, but they cannot conflict with the corporate ones. Second, Review: significant findings from the divisional portfolio — a class of investment that is systematically underperforming, a strategic assumption that is proving wrong — travel upward as input into the corporate portfolio's next Review and its feedback to strategy. Third, Resource contention: scarce resources shared across divisional portfolios — enterprise platforms, shared technical functions, senior executive attention — are managed at the corporate Steer level.

The failure to design the governance connections explicitly produces the standard enterprise pathology: divisions running their portfolios in complete isolation, optimising locally in ways that are incompatible with or destructive to each other, with the corporate level unable to see what is happening below it until something goes wrong at scale.

10.7 How the dials evolve

Organisational capability in portfolio management grows in a recognisable pattern, and the dials should be expected to evolve through it.

At the early stage, the organisation is learning the discipline: making the comparison explicit, maintaining the demand log, running the cadence. The right setting is Lean. Imposing Managed or Enterprise formality on an organisation that has not yet internalised the basic discipline produces paperwork that masks the absence of the discipline underneath.

At the intermediate stage, the organisation has the basic loop running. The demand is growing — more investments, more stakeholders, more complexity than a single owner can hold. The board is needed; the documentation earns its cost. The right setting is Managed, introduced deliberately, with explicit agreement on what the new process requires from everyone involved.

At the mature stage, the organisation's portfolio is large enough, governed formally enough, and consequential enough that the full Enterprise setting — tiered boards, governed processes, formal audit trail — is both necessary and sustainable. The transition from Managed is triggered by scale and stakes, not by ambition.

The common mistake is to rush toward the highest setting. Enterprise governance imposed on a Lean-capable organisation does not make the organisation's portfolio management more rigorous. It makes the organisation's portfolio management theatrical: elaborate processes that nobody uses honestly, because the organisation is not yet capable of running them with the necessary discipline. Theatricality is the opposite of rigor, and it is harder to recover from than honest Lean.

With the operating model set — roles clear, rights explicit, dials calibrated, portfolio structure known — the method has its infrastructure. What it does not yet have is the organisation's willingness to actually live by it. That is the subject of the final chapter.

CHAPTER 11

Making It Stick

Making It Stick

11.1 The method meets the organisation

Everything that has been described in this book is, in a technical sense, straightforward. The logic of the Investment Loop is not difficult to understand. The stages are not obscure. The instruments are practical. The settings are clearly defined. And yet most organisations that attempt to run a portfolio management discipline — that build the reference framework, stand up the board, run the first comparison, establish the cadence — find, within twelve to eighteen months, that the method has eroded. The cadence slips. The comparison becomes a formality. The benefits tracker is maintained but not reviewed. The board approves what arrives rather than comparing it to what else could arrive. The stop decisions are deferred.

This is not because the people are bad at their jobs. It is because applying the method consistently, in a real organisation with real political pressures, real sponsors with real stakes, and real data that is imperfect and sometimes gamed, is harder than understanding it. Making it stick is a different challenge from designing it, and it requires its own thinking.

11.2 Communicating the decisions

The first and most important thing the portfolio can do to make itself durable is communicate its decisions openly: what was funded and why, what was not funded and why, what was stopped and what that means for the people involved.

Opacity protects decisions in the short run. If the criteria are not published, the weights are not shared, and the comparison is not visible, then no individual decision can be individually challenged — because the basis for it is unknown. This feels like a management advantage, and it is not. It is an advantage for protecting bad decisions and indistinguishable from an advantage for protecting good ones. Opacity teaches everyone that the portfolio process is a black box, which teaches them to work around it rather than through it, which is the beginning of the end of any portfolio discipline.

A portfolio that cannot explain its decisions is a portfolio that will not be trusted. A portfolio that is not trusted will be bypassed. A portfolio that is bypassed has ceased to manage the portfolio.

Transparency does not require publishing every number from every scoring session. It requires, at a minimum: which investments were funded and which were not; that the same criteria and weights were applied to both; and what the funded set means for the organisation's strategic priorities, expressed in terms that those priorities' owners can

understand. The communication should happen before people find out through other channels — because they will — and should be direct enough to address the investments that were declined, not just the ones that were approved.

11.3 Carrying people with you

Stopping an investment — or even declining to start one — affects real people. The team that built the business case spent months on it. The sponsor advocated for it and made commitments based on it. The people who were going to work on it had plans. Treating the stop or decline as a purely financial decision, communicated in portfolio language without acknowledgement of the human consequences, produces resistance that is often more durable than the decision that caused it.

Carrying people with you does not mean softening the decision or being unclear about it. It means doing three things consistently. First, explain the decision in terms of the portfolio logic — the forward value, the alternatives, the capacity — not in terms of the investment's quality or the team's capability. A declined investment is not necessarily a bad investment or a bad team; it is an investment whose forward value, at this moment, does not exceed the alternatives competing for the same pool. Second, be specific about what the decision means practically: what happens to the team, the work, the commitments that were made on the basis of the expected approval. Third, do this early, directly, and with the same information given to everyone with a stake in the outcome.

“A decision communicated clearly and early is a decision people can plan around. A decision communicated late, or communicated differently to different stakeholders, is a decision that teaches people the process cannot be relied upon.”

The resistance to stopping work is intense and predictable enough that it is worth treating as a design problem, not a management problem. If the process is designed well — pre-agreed criteria, visible comparisons, decisions communicated openly — the resistance has less to attach to. If the process is opaque, inconsistent, or communicated poorly, the resistance has legitimate grievances alongside the illegitimate ones, and it cannot be distinguished from them. Design the process to deprive the resistance of its legitimate arguments.

11.4 The usual failure modes

Twelve months in, the failure modes look like this.

The cadence slips. The quarterly review becomes semi-annual, then annual, then irregular. Each slip is justified by a specific, local, reasonable reason: the board has too many things on that cycle; a significant programme is at a critical stage; the organisation is in the middle of a restructuring. Each individual slip is defensible. The cumulative effect is a portfolio that is not actually being managed on a rhythm, and where drift — the very condition the method was designed to prevent — has re-established itself between the increasingly infrequent review points. The discipline at this point is not to review whether slipping is sometimes necessary — it is — but to hold the cadence as a default that requires explicit exception, rather than as an aspiration that yields to any pressure.

The comparison becomes a formality. The demand log is maintained, the One-Pagers are submitted, the scoring is done, the comparison table is produced — and then the board approves everything on the list because declining anything would require a difficult conversation. The comparison exists in form and not in effect. The signs are: every investment submitted scores above the funding line; no investment has ever been declined at the comparison stage; the board's discussion of the comparison lasts less than the time it took to prepare it. The fix is to require that the comparison produce at least one decision — at minimum, a funding line that is clearly defined and leaves something below it — and to track whether the decisions actually differ from the submissions.

The benefits tracker is maintained but not reviewed. The template is completed for each investment at go-live. The post-delivery review dates are in the calendar. When those dates arrive, the item is quietly removed from the agenda, or pushed to "next meeting," until it disappears. Eighteen months later, nobody has looked at whether the value arrived. The benefits tracker has become a document-management exercise rather than a learning instrument. The fix is to make the post-delivery review a mandatory agenda item at the board, with an owner accountable for producing the finding.

Status stays green too long. Self-reported status from investment teams systematically lags reality — not usually through dishonesty, but through optimism, through the reluctance to deliver bad news, and through the genuine difficulty of knowing, from inside, whether the problems you are experiencing are temporary or structural. A portfolio that has relied entirely on self-reported status for six months has a portfolio health picture that is approximately as accurate as the investment teams' collective optimism, which is, in most cases, considerably more accurate than it should be. Steer's external view — the dependency map, the resource-contention picture, the escalation log — is the corrective, and it must be maintained actively, not as a backup to be consulted when problems become undeniable.

Sponsors find the workaround. When the portfolio process produces decisions that sponsors disagree with, the instinct is to find a route around it: a budget exception, a direct executive approval, a reframing of the investment as something that doesn't require portfolio governance. These workarounds are individually small and collectively corrosive. Each one signals that the portfolio process is optional for people with enough seniority or leverage, which teaches everyone that seniority and leverage are the real allocation mechanism. The fix is not to punish the sponsor — it is to close the workaround, by ensuring that every route to funding flows through the portfolio process, and by making the single front door a governance commitment, not merely an expectation.

11.5 Data traps

The method depends on data: forward value estimates, cost-to-complete figures, benefit claims, performance data, risk assessments. Most of this data is not known precisely; it is estimated, by people who have reasons to estimate it in particular directions. Understanding the ways in which portfolio data is systematically distorted is part of running the method honestly.

Guessed estimates presented as analysis. A forward value estimate that is built on three assumptions, each uncertain, and presented as a number to two decimal places, is a guess with more significant figures than it deserves. The method does not require precise numbers — it requires honest ones. An estimate of "between five and fifteen million, depending on adoption rate" is more useful than a number like "eight-point-seven million" that implies a precision the analysis cannot support. The portfolio should prefer honest uncertainty to spurious precision, and the comparison process should treat estimates of similar quality as comparable and estimates of widely different quality as incomparable.

Vanity metrics used as value proxies. Measures that are easy to collect and look impressive but do not correspond to the value the portfolio cares about: page views, deployment frequency, satisfaction scores that are never connected to retention, cost-per-transaction without reference to whether the transaction should have happened. These metrics are not without value, but they are not the same as the benefit claims the investment was funded on, and substituting them for those claims produces a portfolio that is measuring its own activity rather than its outcomes.

Gamed status. A portfolio that penalises the bearers of bad news will receive very little bad news. A portfolio that makes it professionally costly to report a delay will see no delays reported until they are undeniable. The design of the Steer and Review processes should make honest reporting the path of least resistance: no-surprises culture means early bad news is welcomed, not punished; the external view from Steer provides an independent check on self-reported status; and the portfolio owner explicitly creates space — in meeting culture, in the framing of reviews — for problems to surface before they become crises.

11.6 Making the rules credible

A reference framework that is sometimes bypassed is not a framework — it is a decoration. A set of stop criteria that is applied to investments whose sponsors lack power, and waived for investments whose sponsors have it, is not a set of criteria — it is a power map that has been given the name of criteria. Rules that are inconsistently applied lose their authority faster than rules that are visibly wrong, because inconsistency destroys the basis on which anyone can trust the process.

The credibility of the rules is the credibility of the portfolio. Rules applied inconsistently do not merely fail to govern the exceptions; they undermine every decision that was made under them.

Making the rules credible requires three things. First, the rules must be visible: the reference framework, the weights, the stop criteria, and the decision-rights framework must be documented and available to everyone who participates in or is governed by them. Second, they must be applied consistently across the full population of investments — the powerful sponsor's investment and the less powerful one, the favourite programme and the unpopular one — with no exceptions that are not themselves governed by a rule. Third, the people responsible for applying them must be accountable for doing so: a portfolio board that routinely overrides the comparison without stated reasons has an accountability gap that the process owner must address.

11.7 Sustaining the rhythm

The hardest thing to sustain is not the process or the tools. It is the commitment to the cadence — the regular moment of re-decision that the method requires. The cadence is what distinguishes a portfolio that is genuinely continuously managed from one that runs a big review once a year and calls it portfolio management. Sustaining it means treating the review as a standing commitment of the people in the governance structure, not as a discretionary activity that yields to schedule pressure. It means keeping the process light enough that the review is not so burdensome that it becomes a reason to skip it, while keeping it rigorous enough that the review produces real decisions.

Over time, the discipline compounds. A portfolio that has consistently run its cadence for two years has two years of comparable data, two years of learning about what the reference framework actually selects, two years of benefits-after-delivery observations. It can see its own patterns, improve its own criteria, and make steadily better decisions. A portfolio that has run an annual review with significant interruptions has none of this; it is making each decision with the memory of the last one and no more.

The promise of this book is not that the method makes portfolio management easy. It is that it makes it possible — consistently, at any scale, and with the discipline that the organisation's resources actually deserve. The method works if it is worked. That is the only condition.

ANNEX A — Direction — Templates

About these templates

The Direction stage produces four instruments: the strategy-on-a-page, the investment criteria and weighting framework (the reference framework), the risk-appetite sheet, and the allocation-buckets framework. The templates below are skeletons — they define the structure and the questions each instrument must answer, at a level of detail proportionate to the Managed setting. Lean adaptations can reduce each to a single page or a set of headings. Enterprise adaptations will add governance fields, version control, and ratification records. A fuller set of worked examples and tool guidance is provided in Book 2.

A1 — Strategy-on-a-page

Purpose: A single-page summary of the organisation's current strategy, expressed in terms the portfolio can use — not as a narrative of ambition, but as a set of specific priorities, success measures, and constraints that will govern the reference framework and its weights.

Complete this template when the strategy is set or refreshed. Review when any of the inputs change.

Field	Content
Period covered	From [date] to [date]
Strategic intent	One paragraph: what the organisation is trying to achieve in this period and why
Strategic priorities	List the 3–5 specific priorities the portfolio should advance (not values or themes — specific objectives)
Success measures	For each priority: how success will be measured, and at what level it becomes meaningful
Constraints	Specific boundaries the portfolio must observe (regulatory, financial, capability)
Risk posture	One sentence: what kind and level of risk the strategy permits in pursuit of the priorities

Field	Content
Owner and review date	Who approved this, and when it will next be reviewed

A2 — Investment criteria and weighting framework (the reference framework)

Purpose: The agreed set of selection criteria and their strategy-set weights. Every investment is judged against this framework; every weight change requires the process in the notes below.*

Criterion	Definition	Weight (%)	Notes on application
Value	The benefit this investment creates — financial, strategic, or operational — measured by [specify the unit the organisation uses], expressed as forward value (value still to come)	[%]	Score 1-5: 5 = clear, high-confidence benefit exceeding [threshold]; 1 = speculative or negligible
Strategic fit	How directly this investment advances a current strategic priority (as listed in A1)	[%]	Score 1-5: 5 = directly named in a priority; 1 = no traceable link
Risk	The likelihood of failure and the magnitude of harm if it fails — to the investment and to the portfolio	[%]	Score 1-5 (inverted): 5 = low probability, low impact; 1 = high probability, high impact
Cost	The total consumption from the pool: money, scarce skills, and management attention	[%]	Score 1-5 (inverted): 5 = low cost relative to value; 1 = high cost relative to value
[Optional additional criterion]	[Define it — if it cannot be defined, do not include it]	[%]	
Total		100 %	

Notes on use:

1. The criteria are the constant; the weights are the variable. The weights are set by the strategy and reviewed only when the strategy changes or when Review provides strong evidence that a criterion is mis-weighted.
1. Criteria must be defined specifically enough that two people applying them to the same investment produce scores within one point of each other. Test this at the first

scoring session.

1. The total weight must sum to 100%. Re-weighting one criterion requires re-balancing the others.
1. Any weight change must be: proposed in writing with the strategic rationale, reviewed by the portfolio board (or owner at Lean), documented here with the date and the approver.

A3 — Risk-appetite sheet

Purpose: A specific, portfolio-level statement of what risk the portfolio will carry in pursuit of return. This governs the balance of the chosen mix — not the risk score of any individual investment.*

Dimension	Statement	Implication for the mix
Overall tolerance	[State the level: e.g. "We will accept up to X% of the portfolio in uncertain-return investments at any one time"]	The maximum proportion of the funded set that may carry high risk ratings
Categories prohibited	[e.g. "We will not fund investments whose primary risk is reputational without executive sign-off"]	Specific risk types that require escalation or are excluded
Run-the-business floor	[e.g. "At least X% of the pool must be allocated to investments that protect existing operations"]	The minimum proportion reserved for operational continuity
Innovation room	[e.g. "We will protect at least X% of the pool for uncertain bets, funded to learn"]	The minimum protected allocation for exploratory investment
Concentration limit	[e.g. "No more than X% of the pool may depend on a single technology, supplier, or strategic assumption"]	The concentration limit applied at the choosing stage

Review trigger: Review this sheet when the strategy changes, when the portfolio's actual risk profile diverges significantly from the stated appetite, or when Review produces evidence that the appetite is set incorrectly.

A4 — Allocation-buckets framework

Purpose: The pre-committed split of the total pool into categories, protecting each category from being crowded out by the others. Complete this if the portfolio is using buckets; otherwise leave blank.*

Bucket	Definition	Approximate allocation (%)	Notes
Run the business	Investments that fund and protect the existing operation	[%]	Includes mandatory / regulatory work
Change the business	Investments that improve and extend what exists	[%]	Moderate certainty of return; medium horizon
Innovation / explore	Uncertain bets funded to learn	[%]	Fund-to-learn structure; high tolerance for failure within the tranche
[Additional bucket if needed]	[Define it]	[%]	
Total		100%	

Rules (mandatory):

- Investments compete within a bucket each period. A strong innovation candidate does not displace a run-the-business investment; it competes with other innovation candidates for the innovation allocation.
- Bucket sizes are resized between periods, by the portfolio board (or owner), aligned to the strategy. A bucket whose size has not been reviewed in more than two periods should be reviewed immediately.
- Buckets protect the category, not the individual investment. A weak investment in the innovation bucket still loses to a stronger innovation investment.
- Reallocation between buckets within a period requires a formal decision, at the board level, with a stated rationale.

ANNEX B — Demand — Templates

About these templates

Demand produces two instruments: the Investment One-Pager and the intake/demand log. Both are structured for the Managed setting; Lean adaptations use a subset of fields; Enterprise adaptations add governance fields and submission windows. Fuller tool guidance is in Book 2.

B1 — The Investment One-Pager

Purpose: The standard framing document that every investment — new or in-flight — must complete before entering the comparison. It states the minimum required to apply the reference framework. It is not a full business case.*

In-flight investments complete sections marked [IF] with current knowledge, not the original case.

Section 1: What and why

Field	Content
Investment name	A specific, unambiguous name (not a project code)
Sponsor	Name and role of the accountable senior individual
Investment lead	Name and role of the person managing delivery
Type	Programme / Project / Product / Initiative / Experiment
Funding status	New proposal / In-flight (year [N] of [N]) / Renewal
What it is	2–3 sentences: what will be done and what will change as a result
Intended beneficiary	Who experiences the value — and who will confirm it arrived
Return claim	The specific, measurable outcome claimed: [value], [unit], [by when] [IF: updated from original]
Confidence level	High / Medium / Low — and the primary reason for the rating

Section 2: Strategic fit

Field	Content
Strategic priority addressed	Which named priority from the strategy-on-a-page (A1) does this directly advance?
Link description	One sentence: the mechanism by which this investment advances the priority
Fit rating	1–5 against the strategic-fit criterion (using the scoring definition in A2)

Section 3: Cost and resource

Field	Content
Total cost to complete	Money remaining to spend [IF: updated cost-to-complete, not original budget]
Cost this period	What the investment requires from the pool in the current funding period
Scarce skills required	Which capabilities are needed that are in short supply across the portfolio
Peak resource period	When the investment's demand on shared resources is highest

Section 4: Risk and confidence

Field	Content
Primary risk	The single most likely cause of failure
Probability	High / Medium / Low
Impact if it fails	On this investment, and on other investments that depend on it
Risk rating	1–5 against the risk criterion (using the scoring definition in A2)
Uncertain bet?	Yes / No — if Yes, this investment is a candidate for fund-to-learn treatment (see Annex D)

Section 5: Forward value (in-flight only)

Field	Content [IF]
Original value claim	What the investment promised when funded
Updated value estimate	What the investment now expects to deliver, based on current knowledge
Cost to complete	Money still to spend — excluding what has already been spent

Field	Content [IF]
Forward-value summary	One sentence: value still to come vs cost still to spend

A completed One-Pager that cannot state a forward value for Section 5 is not ready for re-comparison.

B2 — Intake and demand log

Purpose: The single record of everything competing for the portfolio's pool — new proposals, in-flight investments, and investments flagged for stopping. Maintained continuously; reviewed at every Steer and Review meeting.*

ID	Investment name	Sponsor	Status	Date surfaced	One-Pager complete?	Uncertain bet?	Outcome	Date decided	Notes
[001]	[Name]	[Name]	New / In-flight / Pending stop	[Date]	Yes / No / Draft	Yes / No	Funded / Declined / Deferred / Stopped	[Date]	

Column notes:

- Status: New (not yet in the funded set), In-flight (currently funded), Pending stop (in the stop-or-continue process).
- One-Pager complete? An investment without a completed One-Pager may not enter the comparison.
- Outcome: Every investment that enters the log must leave it with a recorded outcome. An investment that was informally declined and not recorded is an invisible rejection — and invisible data.
- Notes: Record any conditions attached to the decision (e.g. "funded subject to Q2 capacity confirmation") and any subsequent changes.

Maintenance rules:

- Every proposal — from any source, at any stage — enters the log at the point it is formally surfaced, not at the point it is approved.
- Nothing enters the comparison without an entry on the log.
- The log is owned by the portfolio office (or the portfolio owner at Lean). It is visible to all who propose or decide.
- A proposal that is informally discouraged must still be logged and given a formal outcome. Invisible declines are data losses.

ANNEX C — Value & Decide — Templates

About these templates

The Decide stage spans two chapters: How We Value (Chapter 5) and How We Choose the Mix (Chapter 6). Four templates support the full Decide process: the scoring model, the model selector, the portfolio balance and risk-return view, and the capacity-versus-demand view.

C1 — Scoring model

Purpose: Apply the reference framework (A2) to each investment in the pool to produce a comparable weighted score. Run this for every investment — new and in-flight — at each comparison cycle.*

Investment	Value (score / weight)	Strategic fit (score / weight)	Risk (score / weight)	Cost (score / weight)	Weighted total	Forward value	Notes
[Investment name]	[1–5] / [%]	[1–5] / [%]	[1–5] / [%]	[1–5] / [%]	[computed]	[£/€ or qualitative]	[Uncertain bet? In-flight?]

How to complete:

1. Score each criterion 1–5 using the definitions in A2. Apply the same definitions to every investment in the pool.
1. Multiply each score by the criterion's weight (expressed as a decimal: 35% = 0.35).
1. Sum the weighted scores for the total. The total is the basis of comparison; it is not a recommendation — the comparison also requires the capacity view (C4) and the balance view (C3).
1. Record who scored each investment, and the date. Scoring sessions should include at least two scorers to reduce individual bias; significant disagreements should be discussed and resolved before the comparison.
1. For uncertain bets, the value score reflects the option value of the next learning tranche, not a projected full return.

C2 — Prioritisation-model selector

Purpose: Choose the right valuation model for the investment and setting. The scoring model (C1) is the default; this selector identifies when to supplement or replace it.*

Investment characteristic	Recommended model	When to supplement
Heterogeneous portfolio (different types, sizes, sectors)	Weighted scoring (C1)	Add financial appraisal where strong financial data exists
High-flow product or technology queue	Cost of Delay / WSJF	Supplement with weighted scoring for strategic fit
Investment with measurable financial return and defensible estimates	Financial appraisal (ROI / NPV / payback)	Add weighted scoring for strategy alignment
Uncertain bet / experiment	Option value: value the next learning tranche	Do not use financial appraisal — the return is not estimable
Portfolio-level balance view	Value-versus-effort map	Use alongside the scoring model, not instead of it

Model notes:

1. Cost of Delay: Estimate the value lost per week of delay. Divide by duration/effort to get WSJF. Use for sequencing decisions within a category.
1. Financial appraisal: $ROI = (\text{net benefit} / \text{total cost})$. NPV discounts future cash flows at the organisation's cost of capital. Payback = time to recover the investment. Use only where the value is genuinely financial and estimable.
1. Option value: What is the next increment of learning worth? What does it cost? The ratio is the return on the tranche. Record the stop criteria for the tranche.

C3 — Portfolio balance and risk-return view

Purpose: Look at the proposed funded set as a whole and assess whether it is balanced across the four dimensions. Use this at the choosing stage to check the mix, not to score individual investments.*

Dimension	Current proportion	Target range	Gap	Action required
High risk / high return investments	[%]	[%]-[%]	[+/-]	[Reduce / increase / accept]
Medium risk / medium return	[%]	[%]-[%]	[+/-]	
Low risk / low return	[%]	[%]-[%]	[+/-]	
Near-term horizon (< 1 year)	[%]	[%]-[%]	[+/-]	

Dimension	Current proportion	Target range	Gap	Action required
Medium horizon (1–3 years)	[%]	[%]–[%]	[+/-]	
Long / uncertain horizon	[%]	[%]–[%]	[+/-]	
Run-the-business	[%]	Set by bucket or appetite	[+/-]	
Change-the-business	[%]	Set by bucket or appetite	[+/-]	
Innovation / explore	[%]	Set by bucket or appetite	[+/-]	

Concentration check:

1. What shared technology, platform, or supplier dependency do three or more funded investments share?
1. What strategic assumption do three or more funded investments depend on simultaneously?
1. Is any single team, skill, or capability required by more than [threshold – set per organisation] of the funded set?

If any concentration exceeds the appetite limit in A3, note the action required before finalising the funded set.

C4 – Capacity-versus-demand view

Purpose: Confirm that the proposed funded set can be staffed and delivered. The funding line is drawn against capacity, not appetite.*

Scarce capability / resource	Available supply (FTE or units)	Demand from proposed funded set	Surplus / deficit	Resolution
[e.g. Solution architects]	[N]	[N]	[+/-]	[Defer investment X / hire / contract]
[e.g. Data engineering]	[N]	[N]	[+/-]	
[e.g. Executive sponsor time]	[N hours/month]	[N hours/month]	[+/-]	

Rules:

1. Complete this view for every scarce resource identified at the demand stage and by the Steer process.
1. Do not finalise the funded set until the capacity view shows no unresolved deficits.

1. Where a deficit cannot be resolved, reduce the funded set — descope or defer investments — until the deficit clears. Do not proceed on the assumption that capacity will appear.

ANNEX D — Fund — Templates

About these templates

The Fund stage releases money on a rhythm and structures uncertain investments to learn rather than deliver. Three templates: the funding and commitment sheet, the funding cadence calendar, and the staged fund-to-learn card.

D1 — Funding and commitment sheet

Purpose: The record of what the portfolio has committed — to each investment, in what amount, on what conditions, with what review date. The aggregate of all sheets is the current allocation of the pool.*

Field	Content
Investment name	[Name — matches the intake log and One-Pager]
Funding period	[Start date] to [End date]
Funding mode	Rolling / Periodic / Gated / Fund-to-learn
Amount committed this period	[£ / € / \$ — the amount released for this period only]
Conditions	Any conditions attached to the release (e.g. "subject to delivery of milestone X by [date]")
Review date	The date at which this commitment is next examined
Decision rights	Who must approve the next release (portfolio owner / board / governing board)
Total envelope	The maximum total authorised for this investment across all periods
Spent to date	[For in-flight investments: cumulative spend excluding this period]
Cost to complete	[Updated at each review — the money still needed after this period]
Notes	Any material changes since the last review

Rules:

1. A funding commitment sheet is required for every investment in the funded set.
1. The sum of "amount committed this period" across all sheets must not exceed the total pool for the period.

1. The review date is a portfolio commitment, not a suggestion. If the review date passes without a formal decision, the commitment should be treated as lapsed and the investment brought to the next Steer meeting.

D2 — Funding cadence calendar

Purpose: The schedule of the portfolio's regular re-decision points for the coming year — the heartbeat of the loop.*

Date	Event type	Purpose	Preparation due	Decision rights
[Date]	Direction review	Review and update the reference framework and weights	[Date — N weeks before]	Portfolio board
[Date]	Demand window opens	Intake log accepts new proposals		Portfolio office
[Date]	Demand window closes	One-Pagers due for this cycle		Portfolio office
[Date]	Comparison session	Scoring, balance check, capacity check, funding-line decision	[Date — N weeks before]	Portfolio board
[Date]	Funding release	Money committed to the approved set for the period		Portfolio owner
[Date]	Steer forum	Running portfolio health, dependencies, contention, escalations	[Date — N days before]	Portfolio owner
[Date]	Review / re-prioritisation	Re-judge full live set, stop-or-continue decisions, reallocation	[Date — N weeks before]	Portfolio board

Lean adaptation: Collapse the demand, comparison, and review events into a single regular session with the owner. The calendar is the owner's standing commitment.

Enterprise adaptation: Expand to reflect tiered governance cycles, nested sub-portfolio cadences, and the relationship between the portfolio calendar and the corporate budget cycle.

D3 — Staged fund-to-learn card

Purpose: Structure an uncertain investment as a sequence of learning tranches, each funded to answer a specific question. Complete one card per uncertain bet.*

Field	Content
Investment name	[Name]
Classification	Uncertain bet — fund-to-learn
Strategic hypothesis	The assumption this investment is designed to test: "We believe that [X] will [Y], which will [Z]."
Total envelope	The maximum authorised for all tranches combined, if the learning is positive throughout

Tranche structure

Tranche	Learning objective	Method	Amount	Duration	Decision criteria (to proceed to next tranche)	Owner
Tranche 1	What specific question will this tranche answer?	[Prototype / spike / experiment / pilot]	[£]	[N weeks]	Proceed if: [specific, observable condition]. Stop if: [specific, observable condition]	[Name]
Tranche 2			[£]	[N weeks]	Proceed if: [condition]. Stop if: [condition]	[Name]
Tranche 3			[£]	[N weeks]	Final decision: fund to deliver / stop / redesign	[Name]

Rules:

- Each tranche must have a specific learning objective — a question with an observable answer. A tranche with a vague objective ("explore the market") is not a learning tranche; it is unstructured spending with a fund-to-learn label.
- The decision criteria must be specific enough to determine in the room, at the tranche review, whether the condition is met. "Positive user response" is not specific. "At least 40% of test users complete the core flow without prompting" is.
- A tranche review at which the stop condition is met results in a stop — not a renegotiation of the stop condition.
- Stopping at any tranche is a successful outcome, not a failure. The investment did what it was funded to do: it produced information.

ANNEX E — Steer — Templates

About these templates

The Steer stage makes the chosen mix work together by forcing decisions about dependencies, resource contention, sequencing, duplication, and blockers. Four instruments support it: the portfolio dashboard / health view, the dependency map, the decision and escalation log, and the resource-contention view.

E1 — Portfolio dashboard and health view

Purpose: The at-a-glance picture of every running investment — its funding status, health signal, and any active flags. The starting point for every Steer meeting.*

Investment	Sponsor	Funding mode	Period budget	Spend to date	Status	Health	Active flags	Last updated
[Name]	[Name]	Rolling / Periodic / Gated / FtL	[£]	[£]	On track / At risk / Off track / Paused	Green / Amber / Red	[Dependency / Contention / Blocker / Review flagged]	[Date]

Health definitions (apply consistently):

1. Green: Investment is progressing as expected; no material change to forward value or cost-to-complete; no unresolved dependencies affecting the critical path.
1. Amber: A material issue exists that the investment owner is actively managing, with a resolution plan; forward value or cost-to-complete has changed by less than [threshold — set per portfolio]; or a dependency is at risk but has a mitigation.
1. Red: A material issue exists that the investment cannot resolve without portfolio-level intervention; forward value or cost-to-complete has changed by more than [threshold]; or a dependency is blocking progress with no current mitigation.

Rules:

1. Status is set by the portfolio office based on the Steer process — not solely on the investment's self-report.
1. Any investment at Red for two consecutive Steer cycles must be brought to the Review agenda.
1. An investment whose status has been Amber or Green for more than three cycles without any significant change in the underlying picture should be checked for

optimism bias.

E2 — Dependency map

Purpose: A live record of the dependencies between running investments — what each investment is waiting for, from whom, by when, and what happens if the dependency is late.*

Dependent investment	Waiting for	From (investment / team)	Required by	If late: impact	Status	Owner of resolution
[Investment A]	[Specific output / decision / capability]	[Investment B / external]	[Date]	[N weeks delay / blocks milestone / blocks funding gate]	Resolved / On track / At risk / Blocking	[Name]

Critical-path flag: Mark any dependency whose failure would delay the dependent investment's delivery milestone or funding gate with a [CRITICAL] flag. Critical dependencies are reviewed first at every Steer meeting.

Rules:

1. All inter-investment dependencies are recorded here. Bilateral dependency management between teams is permitted but must not substitute for portfolio visibility.
1. A dependency at risk must have a named resolution owner and a resolution date. A dependency with no resolution owner is a blocker waiting to be named.
1. Resolved dependencies are archived, not deleted — they provide a record of the portfolio's dependency history.

E3 — Decision and escalation log

Purpose: The accountability record of the Steer process. Every blocker, portfolio-level decision, and escalation is logged here with an owner and a date. Items are commitments, not observations.*

ID	Date raised	Description	Type	Owner	Resolution required by	Status	Outcome / date resolved
[E001]	[Date]	[Brief description of the blocker, decision, or escalation]	Dependency / Contention / Blocker / Duplication / Escalation	[Name — the person accountable for resolution, not the person who raised it]	[Date]	Open / In progress / Resolved / Escalated	[Outcome and date]

Rules:

1. Every item on the log has a named owner — a specific person, not a team or function — who is accountable for the resolution.
1. An item that is Open with no progress after two Steer cycles is escalated automatically. The escalation owner is the portfolio owner (Lean / Managed) or the governing board (Enterprise).
1. The log is reviewed in full at every Steer meeting. Items are not removed until they are resolved and the resolution is confirmed.
1. The log is the primary input to the Review stage's assessment of Steer's hard truths. The pattern of items on the log — types, ages, resolution rates — is information about the portfolio's health that complements the investment's self-report.

E4 — Resource-contention view

Purpose: Shows where the portfolio's scarce resources are currently allocated, and where competing demands exist. Used at each Steer meeting to make contention visible and to produce resolution decisions.*

Scarce resource	Total available	Allocated to (investment)	Amount allocated	Period	Contested by (investment)	Contest amount	Status	Resolution
[e.g. Solution architects]	[N FTE]	[Investment A]	[N FTE]	[Period]	[Investment B]	[N FTE]	Uncontested / Contested / Resolved	[Investment A gets N FTE; Investment B deferred to Q3]

Rules:

1. Every scarce resource identified at the capacity-check stage (C4) appears here.

1. Contested resources must be resolved at the Steer meeting at which the contest is identified — not carried forward to the next cycle.
1. The resolution is a portfolio-level decision, not a bilateral negotiation. The portfolio makes the call; the teams implement it.
1. A resource that is persistently contested across multiple cycles is a signal that the portfolio is over-committed on that capability. The Review stage should consider whether the funded set should be reduced.

ANNEX F — Review — Templates

About these templates

The Review stage re-judges the live set, tracks benefits after delivery, stops work that no longer earns its place, and feeds learning back to strategy. Three templates: the value and benefits tracker, the stop-or-continue template, and the lessons log.

F1 — Value and benefits tracker

Purpose: Follows each investment from funding through to benefit realisation, confirming whether the value that was claimed actually arrived — and feeding that evidence back into the reference framework and Direction.*

Part 1: In-flight value tracking

Review at each cadence. Updated values are the input to the re-prioritisation comparison.

Investment	Original value claim	Updated value estimate	Confidence	Cost-to-complete (updated)	Forward value summary	Change since last cycle	Action
[Name]	[Original claim from One-Pager]	[Current best estimate]	High / Med / Low	[£]	[Updated value] vs [cost to complete]	Increased / Stable / Decreased / Materially reduced	[Continue / Review / Escalate to stop-or-continue]

Part 2: Post-delivery benefit tracking

Created at go-live. Reviewed at the agreed benefit-realisation checkpoint(s) — typically 6 and 12 months after delivery.

Investment	Delivery date	Original benefit claim	Benefit measure	Checkpoint 1 date	Checkpoint 1 result	Checkpoint 2 date	Checkpoint 2 result	Realised?	Learning
[Name]	[Date]	[What was promised, in measurable terms]	[The metric used]	[Date]	[Actual vs expected]	[Date]	[Actual vs expected]	Yes / Partial / No	[What does this tell us about the criterion or assumption that justified it?]

Rules:

1. Every investment with a benefit claim requires a row in Part 2. The post-delivery tracker is created at the point of go-live — not retrospectively when someone asks whether the benefit arrived.
1. The checkpoint date is set at go-live, based on when the benefit is expected to materialise — not at the point of the review. A benefit expected to arrive twelve months after delivery has a checkpoint at twelve months.
1. "Realised: No" is not a failure of the tracker. It is a finding. The Learning column must be completed: what does this tell us about the investment criteria, the weights, or the assumptions the portfolio has been funding on?
1. Lessons from Part 2 feed the Direction re-weighting conversation. An investment type that consistently fails to realise its claimed benefit is evidence that its scoring is mis-calibrated.

F2 — Stop-or-continue template

Purpose: A structured decision record for any investment that has been flagged for a stop-or-continue decision — whether triggered by the cadence, by a material change in forward value, or by a stop criterion being met.*

Field	Content
Investment name	[Name]
Trigger	Stop criterion met / Material change in forward value / Performance / Periodic review
Original stop criteria	[The specific criteria agreed at the funding stage — from the funding and commitment sheet]
Stop criteria met?	[Yes / No / Partially — specify which criteria and whether they are met]
Updated forward value	Value still to come: [£ or qualitative]. Cost still to spend: [£]

Field	Content
Alternative use of the money	What would the portfolio fund if this investment stops? [Name the candidates — this is the most important field]
Steer input	Summary of the portfolio-level health picture: dependencies, contention, blockers, escalations
Recommendation	Continue / Stop / Conditional continuation (with stated condition and review date)
Decision	Continue / Stop / Conditional
Decision maker	[Portfolio owner / Portfolio board / Governing board]
Date	[Date of decision]
Rationale	One paragraph: why this decision, on this basis
Communication plan	Who is informed, by whom, by when, and in what terms

Rules:

1. Every stop decision is recorded here. An investment that stops without a formal record has created an invisible precedent — visible only to anyone who later wonders why the investment disappeared.
1. The "alternative use of the money" field is mandatory. A stop decision that does not name what the money is reallocated to is a stop without the reallocation — which loses half its value.
1. The rationale must reference the forward value and the stop criteria, not the investment's quality or the team's effort. Stopping an investment is a portfolio decision, not a performance judgement.

F3 — Lessons log

Purpose: A running record of what the portfolio has learned — from stopped investments, from benefit reviews, from Steer findings, and from the re-prioritisation process — that should inform the next Direction conversation.*

ID	Source	Date	Learning	Implication	Action	Status
[L001]	Stop decision / Benefits review / Steer / Re-prioritisation	[Date]	[What was learned: a specific finding, not a generalisation]	[What this implies for the reference framework, the weights, the risk appetite, or the bucket sizes]	[Update criterion definition / Re-weight / Adjust appetite / Change bucket / No action — accept as noise]	Open / Applied / Deferred

Rules:

1. A learning is only a learning if it is specific enough to act on. "We should be more careful with technology risk" is not a learning. "Three investments funded on the assumption that Platform X would scale to N users have failed on that assumption — we have over-weighted the value criterion relative to the technical-risk criterion for platform-dependent investments" is a learning.
1. Open learnings are reviewed at each Direction re-set. An open learning that has been sitting on the log for more than two periods without a decision on whether to act on it is a decision by default — and not the right one.
1. The log is the primary input to the portfolio's feedback to strategy. It is not a report of activity; it is the portfolio's honest account of what its choices revealed about the world.

ANNEX G — Cross-cutting — Templates

About these templates

These four templates apply across the whole loop rather than to a single stage. They support the operating model, the initial dial-setting, and the ongoing communication and governance of the portfolio.

G1 — Roles and decision-rights matrix

Purpose: The definitive record of who does what in the portfolio — their role, their decision rights, and the thresholds at which decisions must be escalated.*

Decision	Portfolio owner / chair	Portfolio board	Governing board / Executive	Portfolio office	Investment sponsor	Notes
New investment approval — below threshold	Decides	Informed	Informed	Prepares	Proposes	Threshold : [define in £]
New investment approval — above threshold	Recommends	Decides	Ratifies if above [£]	Prepares	Proposes	
Continuation / stop — below threshold	Decides	Informed	Informed	Prepares	Informed	
Continuation / stop — above threshold	Recommends	Decides	Ratifies if above [£]	Prepares	Consulted	
Reallocation within period	Decides (below threshold)	Decides (above threshold)	Ratifies above [£]	Logs	Informed	
Bucket resizing	Proposes	Decides	Ratifies	Documents	Informed	
Reference framework re-weighting	Proposes	Decides	Ratifies	Documents	Consulted	

Decision	Portfolio owner / chair	Portfolio board	Governing board / Executive	Portfolio office	Investment sponsor	Notes
Stop criterion override	Cannot override	Can override with stated rationale	Can override	Records decision	May request	
Escalation resolution	Decides below threshold	Decides above threshold	Final escalation	Tracks	Informed	

Complete the threshold values for this portfolio and review annually. A decision-rights matrix with blank thresholds is a matrix without authority.

G2 — The Dial-Setting Worksheet

Purpose: A deliberate, documented choice of the portfolio's setting (Lean / Managed / Enterprise), based on current reality rather than aspiration. Review annually and when the portfolio's scale or context changes materially.*

Step 1: Describe the current reality

Dimension	Current state
Number of investments in the portfolio	[N — include in-flight and planned]
Total pool value	[£ per year]
Number of decision-makers	[Who actively participates in funding decisions]
Reporting relationships	[Who the portfolio answers to; how many governance layers above it]
Data quality	[Honest assessment: are the value estimates reliable? Are benefits tracked?]
Current process maturity	[Is the cadence kept? Are comparisons rigorous? Are stops made?]
Rate of change in environment	[How quickly does strategy or market context shift?]

Step 2: Score the dials

For each dial, assess the current capability honestly — not the desired capability.

Dial	Lean fits if...	Managed fits if...	Enterprise fits if...	Current fit
Cadence	Owner can hold the re-decision in their head; < monthly formally needed	Quarterly structured review is viable and necessary	Formal governance cycles; board-level scheduling required	Lean / Managed / Enterprise
Who decides	One person, or a very small group, can make the calls	A standing board is needed to represent all stakeholders	Tiered boards; nested decision rights by portfolio level	Lean / Managed / Enterprise
Formality	A page and a conversation is genuinely sufficient	A short pack and a meeting is needed	Full documentation, audit trail, and ratification required	Lean / Managed / Enterprise
Funding rigidity	Money can move freely as judgement changes	Periodic cycle governs when money moves	Formal gates; changes require board ratification	Lean / Managed / Enterprise

Step 3: Select the setting

Field	Content
Primary setting	Lean / Managed / Enterprise
Rationale	One paragraph: why this setting fits the current reality
Known tensions	Any dial that is set differently from the primary setting, and why
Review trigger	The condition that would prompt a setting change: [e.g. "Portfolio exceeds N investments" / "Board established" / "Annual review"]
Date	[Date of this assessment]
Owner	[Who approved this setting]

G3 — Stakeholder and communication map

Purpose: Identify every stakeholder group with an interest in the portfolio's decisions, and specify how and when they receive communication about those decisions.*

Stakeholder group	Interest in the portfolio	Information they need	Frequency	Channel	Owner of communication
Executive / governing board	Overall portfolio health; strategic alignment; major stops and starts	Portfolio balance view; funded set summary; significant decisions and their rationale	Each cycle	Board pack	Portfolio owner
Investment sponsors	Their investment's status and funding; portfolio decisions that affect their investment	Funding decisions; stop-or-continue outcomes; reallocation decisions	At each decision point	Direct briefing	Portfolio owner / board chair
Investment teams	Funding certainty for the current period; priorities and sequencing decisions	Funding confirmation; Steer decisions that affect their work	Each period	Portfolio office communication	Portfolio office
Strategy function	What the portfolio is learning about what works	Lessons log; benefits-after-delivery findings; re-weighting rationale	At each Direction review	Structured learning report	Portfolio owner
Business stakeholders	What is being funded and why; what is not funded and why	Funded-set summary; declined investments with rationale	Each cycle	Published decision summary	Portfolio office

Rules:

1. Communication happens before stakeholders hear through other channels. A funding decision communicated after the affected sponsor has already heard informally is a trust damage, not a communication.
1. The rationale for declines and stops is communicated directly and in the same terms as the rationale for approvals: the forward value, the comparison, the criteria. Not a judgement of quality.

G4 — Portfolio health and risk–balance card

Purpose: A summary card of the portfolio's overall health and risk balance, produced at each Review cycle and presented to the governing board. One page.*

Dimension	Metric	Curr ent	Target / appetite	Status	Trend
Total pool committed	[£]	[£]	[£ – total available]	On / Over / Under	Up / Stable / Down
Investments in funded set	Count	[N]	[N target range]	In range / Out of range	
Above-line: new investments	Count and [£]				
Below-line: declined this cycle	Count				
Stopped this cycle	Count and [£] freed				
High-risk investments	% of pool	[%]	[% from appetite A3]	Within / Exceeds appetite	
Run-the-business	% of pool	[%]	[% from A4]	Within / Outside bucket	
Change-the-business	% of pool	[%]	[% from A4]	Within / Outside bucket	
Innovation / explore	% of pool	[%]	[% from A4]	Within / Outside bucket	
Investments with Amber/Red status	Count	[N]	0–[N threshold]	In range / Requires action	
Unresolved escalations > 2 cycles	Count	[N]	0	Resolved / Outstanding	
Benefits realised (last 12 months)	[£ or qualitative]		[Claimed at go-live for same period]	On track / Below	

This card is the portfolio's one-page account of itself. If the card cannot be completed honestly, that is itself a finding about data quality and the portfolio's health.