

Business Strategy as Code: An Auditable Method for Executing a Transformation Strategy

Chapter 4 — The Execution Chain

September 2026

Aliaksei Ivanouski

Abstract

An HX1 pack names choices — an operating model, an enabler grid, a staged roadmap, one falsifiable metric target — and closes there. It does not yet say who is accountable for spending against it, what has been exposed to challenge, what would halt the work, or whether anything has moved. Business Strategy as Code (BSaC_) — the Execution Chain turns an admitted HX1 pack, handed off at the HX2 boundary, into an auditable record of mobilisation, delivery, adoption and realised benefit, never a claim about whether a transformation itself succeeds. It runs eleven Moves, E1 to E11, across four acts. Mobilise names the sponsor, the delivery source and the rollout posture the pack's roadmap left open. Expose pre-mortems the plan and opens a risk register before funding continues. Adopt diagnoses the people the change reaches and builds the coalition and adoption record that follow them. Run & realise operates the cadence that reviews exposure weekly, delivery and adoption monthly, and releases the next wave's funding quarterly — GC3 alone carries that authority, tested on tracked movement of a named metric — closing on a benefits ledger checked against Chapter 1's own success criteria. Three verdicts govern entry: RELEASED, WITHHELD, RETURNED. This paper specifies the chain and marks every rule by level, claiming no more than an execution auditable end to end. No study shows this chain, or naming a sponsor, or running a governance cadence, raises a transformation's success rate — and no defensible base rate for change failure exists in either direction.

About this document

- **Language.** English.
- **Illustrations.** Generic and informative: no company names, no financial figures of any kind, and no fixed count beyond the method's own — four acts, eleven Moves, three verdicts, nine challenges, five failure modes.
- **Nature.** A method description for consultancy and advisory work. It makes no claim of legal, financial, tax, accounting or investment advice, and no rule in this chapter substitutes for it.
- **Status.** Published 22 September 2026; last updated 28 September 2026.
- **Companion guide.** A practitioner guide with detailed procedures, worksheets and a full worked case is in preparation; it is not part of this specification.
- **Relationship to Chapter 3.** This chapter's entry contract admits only an HX1 pack Chapter 3 has already published, handed off at the HX2 boundary; it does not re-open, re-score or re-derive a choice Chapter 3 has made.
- **Relationship to Chapter 1.** This chapter reports movement measured against the success criteria Chapter 1 states; it does not re-diagnose the underlying complication or reopen Chapter 1's own root-cause work.

1. Introduction

1.1 The problem and the use case

A digital-transformation strategy can satisfy every check Chapter 3 publishes — an admitted vision, a chosen operating model, a staged roadmap, a governance design that marks its own funding boundary — and the work it describes can still never begin, or begin and run for a year with nobody able to say what moved. Chapter 3 is explicit about what its own boundary does not reach: mobilising the work, exposing it before anyone outside the wave is told, running the cadence that funds or halts it, and measuring whether the change it promised occurred. A strategy that never crosses into those four questions is a document, not an execution.

This chapter exists to close that gap, and it opens under a constraint that governs everything after it: Hughes's (2011) finding — cited here by finding, not by the number itself — resolved, on inspection, to no published instance carrying valid empirical support (the full citation-check argument sits at §7.5), and is symmetrical: evidence against the widely repeated failure-rate figure, not evidence for a lower, corrected one. There is, at the time of writing, no defensible failure rate for change or transformation programmes in either direction. What follows is a method for making one programme's own claims checkable, one claim at a time, each carrying its own owner, date and closure condition — never a wager on a rate borrowed from elsewhere.

1.2 Conventions

Requirement-level key words (MUST, SHOULD, MAY) follow the shared conventions at whitepaper.strategyascode.org/conventions/. Here they are fixed independently of Chapters 1 through 3, and a MUST binds an input, an output or a record.

Chapter 2 fixes "gate" for its eight quality gates, "stop" for its fifteen numbered units, and "Monday test" for gate 1's consequence check. Chapter 1 fixes "transform" for its own four units. Chapter 3 fixes "Link" and L1 ... L10 for its own ten. This chapter reuses none of the five — where it names one of these words, it is naming the chapter that owns it, never something of its own. It calls its own eleven operations **Moves**, written bare, E1 ... E11, never with a middle dot, which stays a Chapter 2 code-family notation. A code family this chapter mints or inherits that is genuinely a family — a risk register's own entries, for instance — is written $X \cdot 1 \dots X \cdot n$, with $X \cdot i$ naming one instance and n never standing for a count; a Move is never written this way, since a Move is one of eleven fixed integers, not a family.

Each Move admits its input under its own **admission check** and releases its output under its own **contract check**, in exactly the sense the earlier chapters already publish. This chapter's own further vocabulary — the **wave** a Move's output belongs to, the exposure **posture** a wave declares, the **band** a risk entry is scored against, a **register entry**, the **benefits ledger**, and the **independent reviewer** a continuation claim requires — is defined in full where each first arises (§5), not listed here in advance.

The three verdicts this chapter's own entry contract issues are printed in capitals — RELEASED, WITHHELD, RETURNED. GC3's own gate outcomes are printed in lower case — go, conditional-go, recycle, hold, kill. The two sets are different vocabularies over different objects, one chapter's entry decision and another gate's funding decision, and this paper never mixes them. What this chain produces, once every check has passed, is named a **plan** — the first chapter in the series to use the word for its own output. Chapter 3's own product is never retro-named a plan; it is the strategy this chapter's plan executes. This chapter also keeps Chapter 1's word **challenge** for its own adversarial set, with §7.1 stating once how the two relate, following Chapter 3's own precedent.

1.3 How to read this paper

§2 gives this chapter's own foundations. §3 gives the entry contract in full. §4 gives the chain as four acts. §5 gives each of the eleven Moves as a card, then the nine challenges (§5.12) and the Moves-and-rules table (§5.13). §6 covers evidence, provenance and the handoff. §7 covers who may challenge, failure modes and non-claims. §8 covers applicability. §9 sets out the alternatives considered, with its own extended rationale. §10 covers scope, roadmap and references.

1.4 Abbreviations

This chapter mints its own Move and artefact codes (§4); the short forms below are the abbreviations its prose otherwise uses.

- BSaC_ — Business Strategy as Code
- COO — Chief Operating Officer
- CTO — Chief Technology Officer
- ERP — Enterprise Resource Planning
- KPI — Key Performance Indicator
- MBA — Master of Business Administration
- NAO — National Audit Office
- RAID — Risks, Assumptions, Issues, Dependencies
- RFC — Request for Comments
- SRO — Senior Responsible Owner

- VRIN — Valuable, Rare, Inimitable, Non-substitutable

2. Foundations

Chapter 3 draws its foundations from maturity models, digital-domain mapping and business-model literature; this chapter draws on three different bodies of practice instead, and does not restate what the earlier chapters already establish.

Escalation research (Staw, 1976; Brockner, 1992) and its de-escalation counterpart (Keil & Robey, 1999) support the design commitment at §7.2: a continuation decision is more reliable when review sits outside the chain that made the choice; the full argument sits at §9's extended rationale.

Flyvbjerg (2008), Flyvbjerg, Garbuio and Lovo (2009) and Cantarelli, Davis, Pinto and Turner (2025) together support fixing a wave's exit criteria before it starts and reading them against a stated comparison, while leaving the decision itself to a named person, never a threshold.

Conway's (1968) mirroring hypothesis, supported by MacCormack, Baldwin and Rusnak (2012), extends here to a wave's own delivery boundaries: a boundary cut along the wrong seam reproduces a mismatch this literature already names.

The practitioner canon this chapter's vocabulary borrows from supplies naming, never the authority behind a rule, as §6.2 states in full. Where this chapter states a MUST, it rests on the evidence above, on directly verified audit findings, or on a definition this chapter fixes for itself — never on a practitioner framework's own popularity, however widely it is taught.

None of the studies above tested a wave's execution run; each is read here for the shape of a hazard it names, not a validated procedure.

3. The principle

3.1 The rule at the handoff boundary

A pack can satisfy every clause Chapter 3's exit contract publishes and still be unexecutable. Every named component can be present — the vision, the operating model, the portfolio, the governance design, the metric tree — and the pack can still name nobody accountable for spending against it, commit to nothing reversible, and point to no measure anyone will ever actually read. Completeness at the boundary is not readiness to act.

The rule that carried every stop, transform and Link through the earlier chapters — no output leaves without a traceable input behind it — holds at this boundary too, and changes shape once more. What Chapter 1 admitted was a diagnosis; what Chapter 3 admitted was a coded pillar. What arrives here is neither a claim nor a pillar: a decision, already made and already governed, still waiting for a named owner accountable for spending against it. The check that follows does not re-argue anything Chapter 3 already closed; it asks a narrower question — is there now enough on the record for spending, exposing and reversing to be attributed to someone, rather than merely hoped for.

That narrower question has exactly two ways to fail, and this chapter is the first in the series that needs two different words for failing it.

3.2 The entry contract

Chapter 3 publishes what its pack carries: five named components, consolidated under HX1 and crossing at the HX2 boundary (Chapter 3, §5.11). This chapter checks only that all five arrived and that the boundary marker itself names what this chapter owns; a pack short one component is not yet this chapter's input, and a boundary that does not say what crosses it fails the same check for the same reason.

Where this chapter and Chapter 3's exit contract speak to the same requirement, the exit contract governs. This chapter MUST NOT weaken, redefine or restate an HX1 component, the HX2 boundary, or GC3's published purpose. Where it adds a requirement Chapter 3 does not make — a named independent reviewer, the three-verdict release contract below, the two governance tiers beneath GC3 — every one of those additions is strictly stricter than what Chapter 3 published, never looser.


The check issues exactly one of three verdicts, and no other word. RELEASED means every entry condition holds: funding may be drawn against the wave and exposure may begin. WITHHELD means a condition fails on something this chapter can itself supply — an unnamed sponsor, a mandate record citing no funding envelope, an unrecorded separation of approval

from execution, an unnamed authority to call a halt — and nothing is spent or exposed until the named gap is filled, inside this chapter. RETURNED means the failure sits in the pack itself, not in anything this chapter controls: a missing component, or a boundary that never names what this chapter owns. A RETURNED pack is repairable only where it was made, across HX2, back in Chapter 3.

Two failure verdicts exist, not one, because a failure has two genuinely different addresses: its own preconditions, and the pack it was handed. Collapsing them into a single word would send an unnamed sponsor back to Chapter 3, which has none to give, or have this chapter quietly patch a component only Chapter 3 can produce.

All three verdicts are printed in capitals. GC3's own outcomes on a running wave — go, conditional go, recycle, hold, kill — are printed in lower case, and a wave held under one of them is never written HELD: that word was set aside for exactly this collision, so a capitalised entry verdict and a lower-case running outcome are never mistaken for the same kind of word.

Chapter 4



Five component slots — the pack's own fixed set, not a repeating code family.

Figure 3.1 — The entry contract: releasing a pack into execution. A pack is released only once all five components have arrived and the boundary names what this chapter owns; one exit repairs here, one returns across the boundary. This method's format.

3.3 Two-sided contracts, and the one place a second admission check sits

Past this seam, the chain runs on Chapter 3's own pattern, not Chapter 1's: Chapter 3's own Links need only a contract check apiece, because each takes an already-checked upstream artefact as its input. Chapter 4 keeps that shape. One full two-sided check sits at the seam — the entry contract above — and every Move after it carries only its own contract check.

One Move breaks that pattern. E10, the governance cadence, carries a second admission check of its own, because what reaches GC3 is a status claim about a wave still running, assembled by the party accountable for that same wave — the one documented condition under which self-review is known to fail, since a continuation decision taken by the party responsible for the original choice carries the same structural bias whether the reviewer means it to or not. E10's own admission check puts a second, differently positioned reader between that claim and the decision it feeds, precisely where this chapter's opening principle says a second check belongs.

No other Move needs one: each of the rest takes an already-checked output as its input and produces a contract check in turn, the pattern otherwise unbroken.

3.4 What is and is not the contribution

Little in this chapter's own principle is new material — the risk of a self-reviewed continuation decision, criteria fixed before a wave begins, and a delivering structure sharing a shape with what it delivers all predate this paper; §9's extended rationale gives each source in full.

What this chapter adds is the arrangement, not the parts: a three-verdict release contract with two distinct repair addresses, operational review kept separate from GC3's funding authority, an outside reviewer, a banded register, and a benefits ledger that keeps planned apart from realised.

None of this is a claim that a wave released this way performs better, moves faster, or fails less often than one released another way; the claim available here is only that it leaves a record a later reader can audit.

4. The chain

4.1 Four acts at a glance

The chain runs as four acts across its eleven Moves: Mobilise, Expose, Adopt, and Run & realise. Each act closes before the next opens.




Figure 4.1 — The execution chain: four acts, eleven Moves. Read top to bottom: eleven Moves run in sequence through four stacked acts, from the released pack above Mobilise to the benefits ledger below Run & realise. This method's format.

Mobilise (E1–E4) turns a released pack into an accountable, sourced, bounded, postured wave: Mobilisation names a sponsor, mandate and independent reviewer; Sourcing decides, capability by capability, build, buy, borrow or partner, on two independent parents; Delivery boundaries draws work units from the flow of change required, never a reporting line; Rollout posture fixes what is exposed per stage, what cannot be undone, and who may abort it.

Expose (E5–E6) generates and registers what could still go wrong before commitment closes: Pre-mortem treats the failure as already certain and asks why, before the plan locks; Risk register turns each cause into an owned, dated entry whose band carries a declared decision rule.

Adopt (E7–E9) turns "the people side" into three checkable claims: Stakeholder diagnosis names a validated dimension and its observable per group; Sponsorship and coalition binds a person to a named commitment; Adoption measures whether a named behaviour changed in a record, never a message sent.

Run & realise (E10–E11) operates the cadence Chapter 3 designed and closes on Chapter 1's numbers: Governance cadence runs GC1, the weekly exposure review, and GC2, the monthly delivery and adoption review, alongside GC3's stage-gate, which alone releases funding; Benefits realisation separates what was delivered from what changed, entering each benefit against a dated baseline in a planned or realised ledger.

A pack enters this chain only once the entry contract (§3.2) releases it. What leaves is EX1, the benefits ledger, and EX2, the closing boundary naming what is realised, unrealised, owned, and returned as falsified.

4.2 Why this order

The order is not a preference: each act needs what the one before produced, and skipping ahead manufactures an output nobody asked the chain to certify.

Sourcing cannot decide build, buy, borrow or partner before Mobilisation names a mandate and decision rights: nobody can accept its cost until accountability exists. Delivery boundaries cannot be drawn before Sourcing settles what the organisation does itself; a boundary around an undecided capability is a boundary around nothing. Rollout posture depends on both — a stage ladder fixed for a boundary that might still be resourced differently is fixed too early.

Exposure runs after posture: a pre-mortem needs a decision already made to attack, and a cause generated once commitment has closed is a review note, not an output. The risk register needs those causes as one input among several.

Adoption depends on the boundaries and on registered risks whose trigger is a human response: diagnosis is scoped to what they name; sponsorship binds people to Mobilisation's decision rights; its own measures depend on sponsorship's commitments and the pack's leading measures both, since a measure with neither is unfalsifiable.

Folding sourcing into delivery boundaries was considered and rejected: the two cite different parents. Governance cadence runs last because it reviews what every earlier Move produced; benefits realisation closes on a dated post-implementation measure.

4.3 The chain at a glance

The table gives all eleven Moves' shape side by side; §5 gives each in full, as a card.

Move	Act	Takes	Produces	Contract check
E1 Mobilisation	Mobilise	Pack; governance design	Mandate; decision rights; reviewer	Sponsor, mandate, funding envelope, the approval/execution separation and the independent reviewer all recorded
E2 Sourcing	Mobilise	Operating model; capability map	One decision per capability	Differentiation claim or its absence, plus specificity assessment
E3 Delivery boundaries	Mobilise	Sourcing decisions; process dimensions	One boundary per unit	Flow of change served and parent dimension named
E4 Rollout posture	Mobilise	Boundaries; sequenced portfolio	Posture; stage ladder; abort clause	Blast radius, expansion criterion, reversibility, abort authority named
E5 Pre-mortem	Expose	Posture; wave's clearing metric	Candidate causes	Wave and broken artefact named, pre-commitment
E6 Risk register	Expose	Candidate causes; dependencies	Entries; band scale with decision rule	Owner, both dates, pre-set trigger, closure criterion named
E7 Stakeholder diagnosis	Adopt	Boundaries; human-response triggers	One read per group	Validated dimension and its absence-observable named

Move	Act	Takes	Produces	Contract check
E8 Sponsorship and coalition	Adopt	Reads; decision rights	Sponsor, coalition, champion records	One person, one commitment named
E9 Adoption	Adopt	Sponsorship; reads; leading measures	One measure per behaviour	Behaviour, its record, parent measure named
E10 Governance cadence	Run & realise	Register; measures; stage ladder	Exposure, delivery and adoption review; funding stage-gate	One of five outcomes; a go cites measured movement
E11 Benefits realisation	Run & realise	Stage-gate outcomes; metric tree; Ch.1 criteria	Benefit entries; ledger; closing boundary	Dated baseline, owner outside delivery, ledger placement named

Read left to right, a row states what a Move needs before it acts and what it must show once it has; read top to bottom, a Move's Produces is often the next row's Takes. No cell states a value; every field names a code family or an artefact, never a number.


4.4 The three verdicts, and the gate outcomes they are not

The three verdicts governing whether a released pack may begin executing are defined in full at §3.2, and are not restated here. Chapter 1's four verdicts grade a diagnosis's completeness; Chapter 3's two grade a pillar's shape, with nowhere to repair one. This chain decides something neither decides — whether execution may be released against an already-admitted strategy — reusing neither Chapter 3's word for admission nor its word for a referral it does not own.

These three verdicts are not GC3's own outcomes; the two vocabularies are never mixed. GC3 closes a running wave on one of five lower-case outcomes — go, conditional-go, recycle, hold, kill — tested on tracked movement; the entry verdicts test admission once, before a wave exists to run. A capitalised HELD verdict was rejected: beside a lower-case hold outcome it would collide where a reader can least absorb the distinction — why the entry verdict reads WITHHELD instead.

5. The eleven Moves

5.1 E1 — Mobilisation



Purpose. Mobilisation exists to settle, before a wave spends anything or exposes anyone to a change, who is accountable for it, who may approve a change to its scope, budget or schedule, and who stands outside its own chain of accountability with the standing to end it. A pack can satisfy every clause Chapter 3's exit contract publishes and still arrive at an organisation with no answer to any of these three questions, and no name attached to any of them; Mobilisation closes that gap before the first Move that spends or exposes anyone runs.

Inputs. The HX1 pack, released across the HX2 boundary once the entry contract admits it, and Chapter 3's own governance design — the decision rights, the funding model and the review cadence it names for the wave.

Outputs. Three records. MOB1, the mandate record, states the wave's funding envelope and the pack it derives from. MOB2, the decision-rights record, states who approves a change to scope, budget or schedule and who executes the work, kept as two separate names rather than one. MOB3 designates the independent reviewer: a person outside the wave's own chain

of accountability, whose standing to invoke a kill does not depend on the governance cadence's ordinary review cycle.

Derivation rule. MOB1 cites the HX1 pack and the funding envelope Chapter 3's governance design names for it; a mandate with no funding envelope behind it is a statement of intent, not a mandate. MOB2 cites Chapter 3's decision rights and records approval and execution as two separate holders; where one person holds both, the record states that plainly rather than leaving the separation implied. MOB3 names a reviewer whose own position sits outside the wave's reporting line, since a reviewer answering to the sponsor whose own decision they might have to contest carries the sponsor's incentive to continue a failing wave, not an independent one.

Contract check. A named accountable sponsor, a recorded mandate, a funding envelope, a recorded separation of approval from execution and a named independent reviewer all exist. Spend is not authorised while any one is missing.

Anti-pattern. No anti-pattern is named for this Move; the sponsorship failure this chapter guards hardest against belongs three Moves on, at Sponsorship and coalition (§5.8, E8).

Illustration (informative). A programme carries a signed mandate and a funding line but no recorded separation of approval from execution and no reviewer named outside its own chain: the delivery lead approves their own scope changes, and no one outside the reporting line ever held the standing to refuse one.





Figure 5.1 — Mobilisation: the three records. A mandate with no funding envelope behind it, or an approval and an execution holder left as one unrecorded name, does not close mobilisation; the reviewer is designated here, outside the wave's own chain of accountability. This method's format.

5.2 E2 — Sourcing



Purpose. Sourcing decides, capability by capability, whether the organisation builds it, buys it, borrows it or partners for it — and turns that decision into something a reader can check against two independent tests, rather than something merely asserted by whoever holds the budget.

Inputs. MOB2, the decision-rights record; the target operating model's five dimensions and the capability map beneath it, carried in the HX1 pack.

Outputs. SRC · 1 ... SRC · n — one sourcing decision per capability the pack requires, each instance SRC · i naming the capability, the route chosen and the two assessments the route rests on.

Derivation rule. Every sourcing decision cites two parents, tested independently of each other. The first is a differentiation claim: which VRIN property — valuable, rare, imperfectly imitable or non-substitutable, in Barney's (1991) formulation, echoed in the core-competence logic Prahalad and Hamel's (1990) Harvard Business Review article argues for — the capability satisfies, or the explicit, recorded absence of one. The second, tested separately, is an asset-specificity and hold-up assessment, after Tadelis and Williamson's (2012) transaction-cost account: a specific asset invites the party holding it to re-price the relationship once the switch to it is sunk. A decision citing only the first parent is refused, not merely discouraged — routing everything that differentiates to build and everything else to buy mis-routes any activity that does not differentiate but is highly specific, since specificity, not differentiation, is what prices that exposure. A decision resting on unit cost alone tests neither parent and is not recorded.

Contract check. Each sourcing decision cites a differentiation claim, or its stated absence, and an independent specificity assessment; a decision resting on unit cost alone is not recorded.

Anti-pattern. Sourcing by cost alone.

Illustration (informative). A non-differentiating integration is correctly routed away from build on the differentiation test, then correctly routed back towards long-term contracting once the specificity test shows a single supplier could re-price it once the switch to it is sunk.





Figure 5.2 — The sourcing decision's two parents. A decision citing a differentiation claim alone is not recorded: specificity, not differentiation, prices the exposure of everything not built. This method's format.

5.3 E3 — Delivery boundaries



Purpose. Delivery boundaries draws the lines the work will cross — between teams, units or functions — from the flow of change the pack requires, rather than from the reporting lines an organisation already has for other reasons.

Inputs. $SRC \cdot 1 \dots SRC \cdot n$; the operating model's people and process dimensions, carried in the HX1 pack.

Outputs. $DEL \cdot 1 \dots DEL \cdot n$ — one boundary record per team or unit the work will flow across, each instance $DEL \cdot i$ naming the flow of change it serves and the dimension it derives from.

Derivation rule. Conway's (1968) proposition holds that how people coordinate and the shape of what they produce correspond; MacCormack, Baldwin and Rusnak's (2012) matched-pair study of codebases tested that correspondence and found the more loosely coupled organisation's product consistently the more modular of the pair. The chapter states plainly that carrying that finding across from a codebase to a non-technical deliverable is this chapter's own extrapolation, never a re-confirmed result in its own right. Every boundary record therefore names the recurring flow of change it exists to serve and cites the operating-model dimension or capability that flow derives from, rather than citing an existing reporting line as though it were self-justifying. Where a boundary is deliberately drawn to follow an existing reporting line regardless, that choice is recorded as a reasoned exception, not left unstated.

Contract check. Each boundary record names the flow of change it serves and cites its parent dimension; a boundary naming only an existing reporting line is recorded as an exception, or not recorded at all.

Anti-pattern. Org-chart-shaped teams; the label of a coordination framework adopted for its name rather than for the condition it was built to address, with no flow of change ever named as the reason for choosing it.

Illustration (informative). Two units are drawn along an existing approval chain with no flow of change named: every change crosses both, and the coordination cost is paid again on every change that follows rather than designed out once.





Figure 5.3 — Boundaries drawn from flow, not from reporting line. The same units, bounded two ways: a boundary that names no flow of change copies the structure it was drawn from. This method's format.

5.4 E4 — Rollout posture



Purpose. Rollout posture fixes, per wave, what is exposed at each stage, what specifically cannot be undone, and who holds the authority to halt it — so that "phased", "pilot" and "big bang" each name a risk decision rather than a scheduling preference dressed up as one. The posture is the point at which a sequenced portfolio stops reading as a delivery plan and starts reading as a statement of what is exposed, in what order, and on whose authority it can be reversed.

Inputs. DEL · 1 ... DEL · n, the prioritised, sequenced portfolio with its waves and dependencies, carried forward from HX1.


Outputs. ROLL1 the posture; ROLL2 the stage ladder with its blast radius named per stage; ROLL3 the abort clause.

Derivation rule. The posture cites the wave it applies to and states three things: the blast radius exposed at each stage and the criterion for expanding it; the reversibility condition — what specifically cannot be undone, and by when; and the abort criterion, naming the authority who may invoke it, which is never the delivery team alone (Markus, 1983; Hirschheim & Newman, 1988, on resistance tracking the distribution of organisational power rather than design quality alone). Google's own Site Reliability Workbook documents staged-release practice; this chapter's own reading of that practice is that a rollout is only as reversible as its least reversible sub-decision. A full-scope cutover is admissible only where the reversibility condition is actually satisfiable at full scope; otherwise phased or pilot-then-scale is the default, not a fallback reserved for the risk-averse. No comparison of failure rates between postures is stated here: none exists.

Contract check. Each stage names its blast radius and its expansion criterion; the posture names its reversibility condition and its abort authority. A posture naming none of the three is a schedule, not a posture, and is not recorded.

Anti-pattern. d-BANG — the big-bang cutover with no abort path: a full-scope switch with no stated reversibility condition and no named authority to invoke it, whose blast radius is the part of the business the wave exists to protect.


Illustration (informative). A cutover staged across three cohorts, whose least reversible sub-decision — a change to records a later stage cannot cleanly undo — lands in stage one and constrains the whole ladder.



○ = still reversible; ● = the least reversible sub-decision has bound, from stage-k on.
 ⋮ = the range continues; k and n are not counts — stage-k is wherever that binding first falls.

Figure 5.4 — The rollout posture. A posture is only as reversible as its least reversible rung, and the abort lane names who may pull it. This method's format.

5.5 E5 — Pre-mortem



Purpose. The pre-mortem generates the causes of a failure that has not happened yet, before commitment to the plan closes, and hands them to E6 as candidate register entries. It runs ex ante, before commitment, and states the outcome as already certain rather than merely possible — the premise the technique borrows.

Inputs. ROLL1 ... ROLL3; the wave's stated clearing metric, threshold and tested assumption, from HX1.

Outputs. PM·1 ... PM·n — candidate causes, each naming the wave it attacks and the artefact it would break.

Derivation rule. Klein (2007) proposes the technique on the strength of a laboratory finding that treating an outcome as already certain increases the number of causes a person generates for it (Mitchell, Russo & Pennington, 1989). That study measured the count of reasons produced in a laboratory explanation task, not their correctness and not a plan's failure modes; this chapter takes the mechanism and leaves the measurement where it was made. No numeric effectiveness figure is stated for the pre-mortem, anywhere, in any form, and the mechanism is never characterised as surfacing more failure modes than a forward-looking review would.

Contract check. Each candidate cause names the wave it attacks and the artefact it would break, and the session is recorded as having run before commitment closed. A cause generated after the decision is locked is recorded as a review note, not a pre-mortem output.

Anti-pattern. None is named for this Move; the register failure it feeds belongs one Move on, at Risk register (§5.6).

Illustration (informative). A session run the week after sign-off, producing causes nobody can act on because the decisions they attack are already committed.





Figure 5.5 — From candidate cause to register entry. Each cause carries the wave it attacks and the artefact it would break across into the register as a parent an entry can cite; a cause generated after commitment closed crosses as a review note instead. This method's format.

5.6 E6 — Risk register



Purpose. The risk register makes each exposure an owned, dated, falsifiable entry whose band carries an action, rather than a picture that reassures.

Inputs. PM-1 ... PM-n; ROLL1 ... ROLL3; the portfolio's dependencies, from HX1.

Outputs. R-1 ... R-n register entries; RB1 the band scale with its per-band decision rule.

Derivation rule. Every entry cites a parent in an earlier Move or an HX1 component; an entry with no traceable origin is refused. The band scale is ordinal with inline numeric thresholds: a verbal likelihood band carries its numeric range inline at first use in the entry, never in a separate key or a hover, and a severity band is defined by an observable threshold — a cost figure, a schedule slip in weeks, a named contractual breach — never an adjective alone. Wintle et al.'s (2019) study found that agreement with an intended numeric probability band was roughly twice as high when the band's range was printed inline beside the word as when the word stood alone. The register does not rank or prioritise entries by multiplying ordinal likelihood and severity scores: Cox's (2008) analysis shows that this arithmetic has no valid meaning, resolves few hazard pairs, compresses genuine differences in exposure, and for a negatively correlated pair of frequency and severity can rank worse than a random ordering. Cox's critique targets unlabelled multiplication and undefined bands, not ordinal categorisation itself, and this chapter claims no more than that. Each band carries a pre-declared decision rule — the action triggered on entry to that band — which is what converts the scale from a display artefact into a contract. ISO 31000 and

COSO ERM are named here as the register's standards context, never as the source of any matrix design, scoring scale or appetite number; both leave those choices to the adopter. A quantitative estimate may be carried alongside the band where a basis for one exists.

Contract check. Every entry names an individual owner, a creation and a last-reviewed date, an observable trigger written before the fact, and a closure criterion. Closure cites the evidence that satisfies that criterion, never a status update.

Anti-pattern. d-THEATER — the register built once, for a board slide: an artefact indistinguishable in appearance from a live register, its owners named by team, its triggers never written before the fact, never re-reviewed since the wave began.


Illustration (informative). Seven entries, each owned by a team rather than a person, none re-reviewed since the kickoff, and a colour matrix that has not changed while the wave has.

Entry and parent	Owner (individual)	Created, last reviewed	Trigger (before the fact)	Band	Closure criterion
R-1 <i>parent PM-i</i>	<i>named owner</i>	created: date reviewed: date	observable trigger, written before the fact	<i>"likely"</i> <i>+ range</i>	criterion stated, <i>not yet met</i>
⋮	⋮	⋮	⋮	⋮	⋮
R-n <i>parent HX1-c</i>	<i>named owner</i>	created: date reviewed: date	observable trigger, written before the fact	<i>"unlikely"</i> <i>+ range</i>	criterion met — evidence cited

PM, HX1 = an earlier Move's finding or an inbound pack component; i, c = any one item in its set.

Figure 5.6 — The register entry. Every row names a person, two dates, a trigger written before the fact, and what evidence would close it. This method's format.


Band + range (inline)	Observable threshold	Decision rule
B-1 <i>"likely" + range</i>	cost figure	escalate
⋮	⋮	⋮
B-n <i>"unlikely" + range</i>	schedule slip	hold, review



Range and threshold print once, at the band's first use — inline beside the word, never in a separate key. Refused: ordinal likelihood multiplied by ordinal severity has no valid arithmetic meaning.

Figure 5.6a — The bands: ordinal, thresholded, and carrying an action. A band with a range and an action is a contract; the same two axes multiplied together are not. This method's format.

5.7 E7 — Stakeholder diagnosis



Purpose. Stakeholder diagnosis replaces an unfalsifiable label — "there is resistance" — with a named dimension carrying a stated observable, one read per affected group. A read that names only a feeling names nothing a later reviewer could check against the record; a read that names a dimension and an observable can be checked, and can be shown wrong.

Inputs. The delivery-boundary records DEL·1 ... DEL·n; the mandate record MOB1; and the entries on the risk register whose trigger is a human response.

Outputs. SH·1 ... SH·n — one stakeholder read per affected group.

Derivation rule. Every read names which dimension it is asserting. Oreg et al.'s (2008) correlational field study, run across seventeen national samples, found that dispositional resistance to change separates into four measurable factors — routine-seeking, short-term focus, cognitive rigidity and emotional reaction — with an equivalent structure across the samples tested. Holt, Armenakis, Feild and Harris's (2007) psychometric development study found four comparable readiness sub-constructs — discrepancy, efficacy, principal support and personal valence — reporting consistent internal structure across the scale's factors. A read this Move admits names one of these eight factors, never "resistance" or "readiness" as an unqualified whole. A grouping of stakeholders by relative power and interest MAY still be carried, as a working hypothesis rather than a settled classification: it is a heuristic tool with no located predictive-validity study behind it, so a placement made under it carries a stated re-evaluation trigger rather than being fixed once at programme start and left unvisited.

Contract check. Each read MUST name a validated dimension and a stated observable that would show it absent; a read asserting resistance or readiness without naming a dimension is not recorded.

Anti-pattern. None is named at this Move; the comms-only failure the diagnosis exists to prevent is named at E9, where a message answering an unnamed dimension is the shape the failure takes.

Illustration (informative). One group's read names a short-term-focus factor and states what evidence would show it absent; a second group's read says only "resistant to the change" and so names nothing anyone could later confirm or refute.





Figure 5.7 — The stakeholder read: a named dimension and its observable. One read per affected group, each naming a factor from a validated instrument and what would show that factor absent; a placement by relative power and interest is carried beside the read as a working hypothesis, never as the read itself. This method's format.

5.8 E8 — Sponsorship and coalition



Purpose. Sponsorship and coalition bind named people to named decisions and resource commitments, so that "sponsorship" resolves to something a later reader can check against a record, rather than a claim resting on a person's visibility at a launch event.

Inputs. SH · 1 ... SH · n; MOB1; MOB2.

Outputs. CH1 the sponsor commitment record · CH2 the coalition record · CH3 the champion network.

Derivation rule. Every sponsor entry names one specific person and one specific decision or resource commitment that person, and no one else, can make, cited to the decision-rights record MOB2 carries. A role title, a standing committee or a signature on a charter is not a sponsor entry: none of the three names a person who could later be asked whether the decision was actually theirs to make. This is a definitional requirement, not an evidentiary one — a diffuse claim of "leadership support" cannot be falsified, because it names no one accountable for a specific act. Appelbaum, Habashy, Malo and Shafiq's (2012) review of the empirical support behind Kotter's change model argues that support for a "guiding coalition" holds unevenly step by step, with no study testing the coalition step in isolation or the sequence as a whole — a critique of that model's own evidentiary basis, offered here as rationale for naming individuals, not as the authority the rule itself rests on. Where a change decision is claimed to have been made fairly, the record discloses which of engagement,

explanation and expectation-clarity were present for the affected group, and which were absent. Kim and Mauborgne's (2003) Harvard Business Review article frames these three components as "fair process" — a practitioner synthesis, not a measured finding — and the weight behind disclosing them rests on the underlying organisational-justice research tradition the article draws on, never on the synthesis itself. No vendor benchmarking figure about sponsorship appears in this chapter, in any form.

Contract check. Each sponsor entry MUST name one person and one decision or resource commitment that person can make; an entry naming a role, a committee or a signature is not recorded.


Anti-pattern. *d-SPONSOR — the checkbox sponsor.* An executive who appears at the launch and does not appear again, recorded as a tick rather than as a person holding a decision only they can make.

Illustration (informative). A launch deck whose sponsorship line is a single tick beside an executive's name, against a record naming that executive and the one budget reallocation only they are authorised to make.



Figure 5.8 — The sponsorship cascade: named people, named decisions. Every entry binds one person to one decision or resource commitment that person can make, cited to the decision-rights record; a role, a committee or a signature binds nobody. This method's format.

5.9 E9 — Adoption



Purpose. Adoption measures whether behaviour changed, in a named record, rather than whether a message was sent, a session was attended or a survey was completed.

Inputs. CH1 ... CH3; SH.1 ... SH.n; the metric tree's leading measures, carried forward from HX1.

Outputs. ADO.1 ... ADO.n — one adoption measure per behaviour the pack requires to change.

Derivation rule. Every adoption measure names the behaviour it tracks, the record the measurement is read from, and the HX1 metric or capability it serves. An adoption claim resting on training attendance, awareness-survey completion or communication reach alone is refused: each of those measures exposure to a change, never a change in what anyone actually does. Where a feedback channel is claimed to be genuine rather than nominal, the claim cites at least one recorded instance of dissent that changed a decision. Edmondson's (1999) field study of workgroups found psychological safety to be a team-level condition, distinguishable across teams and linked to observed learning behaviour — a channel's presence on an org chart is not evidence that the condition it depends on is present. Bridges' "neutral zone" and ADKAR's reinforcement stage are named once here, as vocabulary for the interval between an old practice ending and a new one holding steady, and carry no rule of their own.

Contract check. Each adoption measure **MUST** name the behaviour, the record it is read from, and its parent measure; attendance, reach and survey completion are not admissible as adoption evidence.


Anti-pattern. *d-COMMS — change as communication only.* A plan of messages, sessions and a channel: not wrong, and nowhere near sufficient, because nothing in it would read differently whether or not anyone changed what they did.

Illustration (informative). A change plan built from a newsletter, two briefings and a channel — each real, none a record that would ever read differently whether or not anyone changed what they did.

Behaviour	Record it is read from	Parent measure
ADO · 1 <i>behaviour named</i>	<i>record named</i>	HX1 <i>leading measure</i>
ADO · 2 <i>behaviour named</i>	<i>record named</i>	HX1 <i>leading measure</i>
⋮	⋮	⋮
ADO · n <i>behaviour named</i>	<i>record named</i>	HX1 <i>leading measure</i>
Refused as adoption evidence <i>measures exposure, not behaviour</i>		
~~Attendance~~	~~Reach~~	~~Survey completion~~

Figure 5.9 — The adoption ledger. Each row names the behaviour and the record it is read from; the footer holds the three classes of evidence this Move refuses. This method's format. Source file:

5.10 E10 — Governance cadence



Purpose. To run the cadence Chapter 3 designed: the weekly exposure review and the monthly delivery and adoption review, neither of which can touch funding, and a periodic gate that can.

Inputs. R·1 ... R·n and ADO·1 ... ADO·n, the risk register and adoption ledger; ROLL2, the stage ladder with its blast radius named per stage; MOB2 and MOB3, the decision-rights record and the independent-reviewer designation; Chapter 3's governance design and each wave's clearing metric and threshold.

Outputs. GC1, the weekly exposure review; GC2, the monthly delivery and adoption review; GC3, the quarterly funding-continuation gate, inherited from Chapter 3.

Admission check. The wave's continuation criteria exist as recorded before the wave began, and the status claim cites a dated measurement. A claim carrying neither is refused entry to the gate.

Derivation rule. GC1 and GC2 carry no funding authority. GC1 reviews the risk register and adoption ledger; GC2 reviews wave scope, schedule and dependencies. Either may require GC3 to convene ahead of its scheduled date, but neither releases, withholds or withdraws funding. GC3 alone is funding-bearing: Chapter 3 publishes that GC3 governs whether a wave's funding continues, tested on tracked movement, releasing each wave's budget only against K0 movement. This chapter fixes only GC3's mechanics — who convenes it, its closed outcome set, what a go must cite, and who may invoke a kill.

A wave's continuation criteria are recorded before the wave starts and are not edited at the gate that tests them: a threshold negotiated at the review is set by whoever is most invested in a go. The authority to invoke a kill is the independent reviewer named at E1, never the wave's own sponsor or delivery lead. The justification sits in the escalation-of-commitment literature: a decision-maker responsible for a prior choice tends to increase commitment to it rather than reverse it (Staw, 1976; Brockner, 1992), and Keil and Robey's (1999) exploratory review of documented de-escalation cases found that, in most of the cases studied, a failing initiative was stopped by an actor outside the project's own management rather than by that management itself. A hold carries a resolution deadline after which it converts to kill absent a new go; a conditional-go names its condition and the date by which it converts; a recycle records its rework list and re-test date at the same gate.

Contract check. Every gate outcome is one of go, conditional-go, recycle, hold or kill; a go cites the measured movement — metric, value, date — it rests on. A narrative summary is never recorded as a go.

Anti-pattern. The gate that only ever says go; status by assertion.

Illustration (informative). A gate convened on the schedule of the party under review, chaired by the person who chose the wave, closing on a summary nobody can trace to a number.





Figure 5.10 — The cadence: two tiers that cannot fund, one that can. The weekly exposure review and the monthly delivery and adoption review escalate and convene; only the funding gate releases, and the kill path runs outside the wave's own chain. This method's format.

5.11 E11 — Benefits realisation



Purpose. To separate what was delivered from what changed, and to close the chain on the numbers Chapter 1 stated rather than on the ones that happened to move.

Inputs. GC3 outcomes; ADO·1 ... ADO·n, the adoption ledger; the metric tree and North Star, from Chapter 3; Chapter 1's stated success criteria — baseline, target, date.

Outputs. B·1 ... B·n benefit entries in two ledgers, planned and realised; EX1, the benefits ledger; EX2, the closing boundary.

Derivation rule. A benefit entry cites a dated pre-change baseline measurement and names one accountable owner, distinct from the delivery team. An output — a thing the wave delivered — is never recorded as a benefit. A benefit is never marked realised because an output reached go-live; realisation requires a dated post-implementation measurement against the stated baseline. The UK National Audit Office's review of the final benefits statement for the NHS national programme

found roughly two-thirds of forecast benefits still classified as future benefits, and 98% for three constituent programmes. An entry whose only evidence is a business-case estimate with no scheduled confirmation is admitted to the planned ledger alone.

Where two or more initiatives could claim the same underlying movement, an attribution rule is recorded before either reports it as realised. Where independently sourced data exists outside the programme, the realisation status cites it rather than resting on the delivery team's own report — a safeguard against the optimism bias in benefit estimation that government auditors have separately documented (National Audit Office, *Over-Optimism in Government Projects*, 2013). A benefit may be non-financial, provided it still carries a baseline and an owner. No claim is made that benefits-management practice causes success: Serra and Kunc's (2015) cross-country study reports only an association with value-creation activity and a weak, inconsistent relationship with overall perceived success, noting that no empirical evidence of the practice's effectiveness existed before it.

Contract check. A benefit entry names a dated pre-change baseline and one accountable owner outside delivery, and sits in the planned or the realised ledger. An output, or a benefit claimed at go-live, is not admitted to the realised ledger.

Anti-pattern. "We shipped it, so it is done" — a closeout that records delivery and calls it a result.

Illustration (informative). A closeout recording eight delivered initiatives and no re-measured baseline: complete on everything the programme controlled, silent on the only number it was commissioned to move.




Figure 5.11 — The benefits ledger: planned and realised. A benefit crosses on a dated post-implementation measurement, never on a go-live date. This method's format.

5.12 The nine challenges

1. Authority. "Who may halt this, and are they outside the chain that wants it to continue?" Pass condition, MUST: a named person may invoke a kill outcome, and that person is neither the wave's sponsor nor its delivery lead. The escalation research behind this challenge sits at §2 and, in full, at §9's extended rationale. A record naming only insiders as able to invoke a kill has not answered this challenge. Failing it fails E1's mobilisation record, or reopens E10's own governance review, because an authority the record cannot name is an authority the record cannot exercise.

2. Reversibility. "Name the one thing this rollout cannot undo, and by when." Pass condition, MUST: the posture names a reversibility condition and the stage at which it binds. Markus's (1983) account of introducing new technology as a power relation, not a neutral technical event, is the reason a posture silent on reversibility reads as consensus rather than as an unexamined one-way door. Staged-release practice supplies the pattern the posture borrows: expand only past a named, checkable point. A ladder with no rung marked as the least-reversible one has not stated a posture; it has stated a schedule with a date attached. Failing this challenge fails E4 outright.

3. Pre-commitment. *"Show me the continuation criteria as they were recorded before the wave started."* Pass condition, MUST: the criteria exist, dated before the wave, and were not edited at the gate outcome that tests them. Flyvbjerg (2008) and Lovo, Cristofaro & Flyvbjerg (2023) treat criteria set after work has begun as a symptom of optimism bias and strategic misrepresentation, not as prudent flexibility. A continuation case argued from criteria drafted at the review itself has produced a justification, not a test. Failing this challenge fails the status claim E10 carries into GC1 or GC2, whichever cadence tier is reviewing it, and the wave's funding continues unreviewed.

4. Status evidence. *"What observable measurement supports that status, and from what record?"* Pass condition, MUST: the continuation claim cites a metric, a value and a date. Chapter 3's own published standard for GC3 — funding released only against tracked movement of a named metric — is the bar this challenge extends downward to GC1 and GC2, so an unevicenced status claim cannot simply wait for the funding tier to catch it. A status recorded as "on track" with no cited measurement is an opinion wearing a record's shape. Failing this challenge fails E10's own review, whichever cadence tier receives it, before the claim reaches GC3.

5. Band discipline. *"Read me the numeric range behind that likelihood word, and the action it triggers."* Pass condition, MUST: the band carries an inline numeric range and a pre-declared decision rule, and no entry is ranked by multiplying ordinal scores against each other. Cox's (2008) critique of ordinal risk matrices is the reason: two scores multiplied together produce an arithmetic that looks precise and is not, and Wintle et al. (2019) show verbal likelihood terms read inconsistently across people without a stated number attached. A band with a word and no range has not passed. Failing this challenge fails E6's register entry.

6. Sourcing exposure. *"What is the specificity of this relationship, and what does it cost to switch?"* Pass condition, MUST: the decision cites an independent specificity and hold-up assessment, not a differentiation claim standing alone. Transaction-cost reasoning (Tadelis & Williamson, 2012) treats specificity and switching cost as the exposure a sourcing choice actually carries, separately from whatever case is made for why a capability differentiates the wave. A build/buy/borrow/partner decision argued only from differentiation has priced one side of the choice and left the other unexamined. Failing this challenge fails E2.

7. Adoption evidence. *"Show me behaviour in a record, not attendance."* Pass condition, MUST: at least one adoption measure names a behaviour and the record it is read from. This follows from what an adoption measure is by definition: attendance, reach and completed surveys record contact with a change, not a change in what anyone does. Where the behaviour in question is dissent rather than compliance, Edmondson's (1999) psychological-safety research supplies the variant: the record must show that disagreement was voiced, not merely permitted. Failing this challenge fails E9.

8. Resistance naming. *"Which named dimension of readiness or resistance is this, and what would show it absent?"* Pass condition, SHOULD: each read names a validated dimension and a stated observable. Oreg et al.'s (2008) cross-national measure and Holt et al.'s (2007) readiness scale both decompose resistance into named, distinct dimensions rather than one undifferentiated mood in the room. A stakeholder read that reports "resistant" without naming which dimension, or what absence of it would look like, has described a feeling, not an assessed property. Departing from this challenge is recorded, not silently skipped. It bears on E7.

9. Benefit attribution. *"What would have happened without this, and who else is claiming the same movement?"* Pass condition, SHOULD: the benefit states a counterfactual assumption, and any competing claim carries a recorded attribution rule. Zwikael & Smyrk (2012) ground the counterfactual requirement; double-counted benefits, claimed by more than one

initiative against the same movement, are a documented pattern in practitioner portfolio-management writing, not a rare accident. A ledger entry with no stated counterfactual and no attribution rule cannot be told apart from a coincidence. This challenge bears on E11.

CHALLENGE	LEVEL	MOVE IT CAN FAIL
1 Authority	MUST	E1 · E10
2 Reversibility	MUST	E4
3 Pre-commitment	MUST	E10
4 Status evidence	MUST	E10
5 Band discipline	MUST	E6
6 Sourcing exposure	MUST	E2
7 Adoption evidence	MUST	E9
8 Resistance naming	SHOULD	E7
9 Benefit attribution	SHOULD	E11

*Nine challenges; MUST challenges are run and recorded whatever the outcome.
A Move named here is the one that challenge tests; it names no verdict and no gate outcome.*

Figure 5.12 — The nine challenges, and the Moves they can fail. Each challenge carries its own requirement level and can send back the output of one or more named Moves; the figure draws that mapping and nothing beyond it. This method's format.

5.13 The Moves and their rules

Each Move's admitted input, produced output and contract check are given once, in full, at §4.3; this appendix does not repeat them. One row per Move instead states its rules by content, never by internal identifier, each marked by level. Every Move runs inside the entry contract on Chapter 3's HX1 pack, admitted at HX2, and the EX1→EX2 exit contract that closes the benefits ledger.

Move	Key rules
E1 Mobilisation	The mandate record MUST name the accountable sponsor as an individual and MUST cite the funding envelope Chapter 3's governance design names; a record naming neither is not a mandate record. Approval authority over scope, budget and schedule MUST be named separately from execution. An outside reviewer MUST be named before the first review.
E2 Sourcing	A sourcing decision MUST cite a differentiation claim naming its property, or its absence, and an independent hold-up assessment. A build lacking a qualifying property MUST be re-routed to buy, borrow or partner. Unit cost alone MUST NOT justify the decision.
E3 Delivery boundaries	A team or unit boundary MUST name the recurring flow it serves and cite the operating-model dimension it derives from. One following an existing reporting line SHOULD be recorded as a reasoned exception.
E4 Rollout posture	A rollout posture MUST state its reversibility condition — what cannot be undone, by when — the exposure at each stage and its expansion criterion, its abort criterion, and who may invoke it, never the delivery team. A full-scope cutover MUST NOT be chosen where reversibility cannot be met at full scope.
E5 Pre-mortem	A pre-mortem SHOULD run before commitment closes; a cause generated afterward is a review note, not a pre-mortem output. No numeric effectiveness figure for it is stated, and hindsight is never characterised as surfacing more failure modes.
E6 Risk register	Every entry MUST name an owner, creation and last-reviewed dates, an advance trigger and a closure criterion, and MUST cite a parent Move or the HX1 pack, else be refused. A likelihood band MUST carry its numeric range inline; severity MUST rest on an observable threshold. Entries MUST NOT be ranked by multiplying likelihood and severity; each band MUST carry a pre-declared action, and closure MUST cite satisfying evidence.

Move	Key rules
E7 Stakeholder diagnosis	A resistance or readiness entry MUST name which validated dimension it asserts and what would show it absent. A power-and-interest placement MAY be a working hypothesis but MUST carry a stated re-evaluation trigger, not a fixed reading.
E8 Sponsorship and coalition	A sponsor entry MUST name a specific person and a specific decision or resource commitment they can make; a title, committee or signature is not one. A fairly-made decision's record SHOULD disclose which of engagement, explanation and expectation-clarity were present, and which absent.
E9 Adoption	An adoption claim MUST cite a behavioural measurement naming the record it is read from; attendance, survey completion and communication reach MUST NOT count as evidence. A genuine feedback channel SHOULD cite a recorded instance of dissent that changed a decision.
E10 Governance cadence	GC1 and GC2 MUST NOT release, withhold or withdraw funding; GC3 alone is funding-bearing, in one of five named outcomes. A conditional-go MUST name its condition and its conversion date, a hold MUST carry a deadline after which it converts to kill, and a recycle MUST record its rework list and its re-test date at the same gate. A continuation decision MUST cite the measured movement it rests on, recorded before the wave starts and left unedited at the review. The kill authority MUST be named in advance, never the wave's sponsor or delivery lead.
E11 Benefits realisation	Every declared benefit MUST cite a dated pre-change baseline and name one owner distinct from delivery; an output MUST NOT be recorded as a benefit. Realisation MUST require a dated post-implementation measurement against that baseline — go-live is not realisation — else the benefit sits in the planned ledger alone. Where initiatives could claim the same movement, an attribution rule MUST be recorded before either reports realisation.

Every Move's output cites a parent reaching the HX1 pack or a Chapter 1 success criterion; the ledger closes at the EX1→EX2 boundary alone.

6. Evidence, provenance and the handoff

6.1 What the execution graph shows


A Move's output earns its place in the record exactly as an admitted pillar earned its own in Chapter 3: nothing enters the chain uncited. Every artefact a Move produces cites a parent, and a parent is always one of three things — a named HX1 component that arrived at the boundary, an earlier Move's own output, or, only at the chain's far end, one of Chapter 1's stated success criteria.

A benefit entry shows the trace at its full length: it cites GC3's outcome and the adoption measure beneath it; the adoption measure cites the coalition record and the stakeholder read beneath that; the stakeholder read cites the boundary a change crosses and the mandate that approved it; the boundary cites the sourcing decision it depends on; the decision cites the decision-rights record the chain's first Move produces. Follow any benefit entry back this way and the trail ends at the HX1 pack, never at an unwritten assumption.

This chapter carries no filled register of its own. Chapter 3's Appendix B walks a full derivation instance because nothing in it needs a named person, a date or a live figure to stay generic; a running wave cannot be shown the same way without becoming exactly the named case this chapter refuses (§7.5). The graph's rule is stated once, at every Move's own card, and holds by construction rather than by a worked example.

Where a Move's own derivation states a general claim rather than naming a specific parent outright, the card says so at the point the claim is made, on the same discipline Chapter 3 states for its own inferred edges (Chapter 3, §6.3) — never


presented as carrying the weight of a citation the Move states directly.



— solid = a stated parent — highlighted = a different relation [] muted = no path forward
i = any one code in its family; this chapter registers the rule, not a worked instance.

Figure 6.1 — The execution graph, schematic (code families only). Every Move's output cites a parent — an HX1 component, an earlier Move's output, or, at the chain's far end, one of Chapter 1's stated success criteria; this chapter

registers the rule rather than a worked instance. This method's format.



Arrows point from a code to the parent it cites; dashed arrows lead to a real but untraced parent.

This is one path through a graph where some codes cite more than one parent — not the whole graph.

Figure 6.2 — Reading a benefit entry like a stack trace: one path back to the pack (informative). Start at a benefit entry and follow each code's cited parent down until the trail reaches the HX1 pack; where a code cites more than one parent the figure follows one of them, and the others are as real. This method's format.

6.2 The evidence bar this chapter applies to itself

A rule at MUST level in this chapter rests on one of three things only: a peer-reviewed empirical finding, a directly verified institutional audit's own published conclusion, or a distinction that needs no measurement to hold because it is true by definition. The separation this chapter draws between an output a programme delivers and a benefit it realises is the last kind: delivering an output is not the claim that anyone's work changed because of it, and the two need no study to stay apart.

Kotter's steps, ADKAR, Prosci's own model, Bridges' transitions, Lewin's stages — a triad no primary Lewin text states in that form (Cummings, Bridgman & Brown, 2016) — MSP, RAID, Mendelow's grid, Team Topologies, and the Spotify/SAFE operating patterns appear throughout this chapter for a different reason entirely: each supplies a name a reader already carries for a shape or a failure mode, never the evidence a rule stands on. So does a borrowed phrase such as "watermelon reporting". None of this vocabulary carries a MUST anywhere in this chapter, and a reader is told so at first use rather than left to infer which register a sentence belongs to from its tone alone.

A vendor's own benchmarking survey sits in neither class, however consistent its own figure looks across editions. It cannot carry a MUST: no independent hand has checked a number a vendor collects from its own client base and reports in its own report. Nor does it sit safely alongside a named framework the way the practitioner canon above does — a framework is at

least a stable idea a reader can go and read in full; a benchmarking figure is a number tied to a method of collection this chapter cannot audit. Where such a figure appears at all, it names the vendor and the word "survey" in the same sentence, and states plainly that the number is self-reported.

6.3 Tracking in flight: from the metric tree to movement

Chapter 3 states the North Star and the tree of metrics beneath it; this chapter reads that tree in flight and records what actually moved. It mints no metric family of its own. Where execution needs a measure the tree does not carry — an adoption measure naming a changed behaviour, a register entry's own trigger — that measure is recorded as this chapter's own execution measure and is never presented as though it belonged to Chapter 3's tree, however closely the two move together in practice.

The distinction matters most exactly where movement disagrees with what a pillar assumed. Chapter 3 publishes that an admitted pillar whose assumption is later falsified is referred back, formally, to Chapter 2's stop 15, and that Chapter 3 itself **MUST NOT** repair the pillar in-chain. Reading the tree in flight is precisely where such a falsification first becomes visible — a leading measure this chapter carries stops moving the way the pillar assumed, or moves the wrong way entirely. This chapter routes that finding back through the same published referral Chapter 3 already states; it **MUST NOT** open a second channel into Chapter 2, and it **MUST NOT** repair a pillar, a vision or an operating model itself, however directly the finding was produced.

One boundary, in a single sentence: a metric family belongs to whichever chapter minted it, and only Chapter 2's own stop reopens what Chapter 3 already closed. Everything this chapter reads in flight is evidence for that referral, never a substitute for it.

6.4 The exit contract: from the pack to the ledger

This chapter's own output earns a code of its own, in two parts. EX1 consolidates the benefits ledger's planned and realised entries into the single record that closes the chain; EX2 marks the boundary it crosses, naming three things a reader can check without being in the room: what is realised against its dated baseline, what is unrealised and still owned by a named person, and what is returned as falsified, rather than quietly dropped.

The HX1 pack was the only connection Chapter 3 opened into this chapter; the re-measured version of Chapter 1's own success criteria is the only thing this chapter hands back. There is no side channel in either direction, which is why a falsified benefit travels the referral §6.3 already names, rather than a line this chapter invents for the occasion.

A benefit crossing into the realised ledger without a dated baseline, or without an owner outside the delivery team that built it, has not crossed at all — it stays in the planned ledger, unresolved, however far past go-live the calendar has moved. The naming discipline is the one Chapter 3 keeps at its own boundary: what closes this chain is named the benefits ledger everywhere this paper states it, never a status report and never a delivery summary, because the two are not the same act.




Figure 6.4 — The exit contract: from the pack to the ledger. EX1 consolidates what was measured; EX2 marks the line it crosses, naming what is realised, what is still owned and what returns falsified. This method's format.

7. Challenges, failure modes and non-claims

7.1 The bar

A Move's output can satisfy its own contract check in full and still be the wrong decision to carry forward — the check tests what the record must contain, never whether the choice behind it should stand. A posture that names a reversibility condition and the stage at which it binds, a register entry that carries an owner and a trigger, a status claim that cites a metric and a date: each can be complete on the record and still be wrong in a way its own check cannot see, because the check was built to test presence, not judgement. The bar this section sets is therefore adversarial, not procedural: a Move's output is admitted only once it has been put in front of someone equipped to break it, and has survived.

The nine challenges (§5.12), fixed as this chapter stands, are numbered separately from Chapter 1's nine and Chapter 3's eight. The shared shape — a fixed, numbered adversarial test set with pass conditions and named verdicts — reflects one shared design decision across the paper, following Chapter 3's own precedent for the same word: never a shared roster, and never a constraint any of the three chapters places on the others. Every challenge **MUST** be run and recorded, whatever the outcome; an unanswered challenge **MUST NOT** be recorded as withstood. A wave that has withstood all nine is challengeable, not thereby right to continue: it has been put in front of people equipped to halt it, and was not halted.

7.2 Who may challenge

Competence to challenge is fixed by what a Move's output claims, not by seniority. A register entry that names an owner and a trigger is challenged most sharply by whoever holds the threshold the trigger names, not by whoever chairs the wave's steering meeting. A posture that names a reversibility condition is challenged, in addition, by whoever would need to invoke the abort authority it names. This mirrors the challenger model the earlier chapters already establish, applied here to a Move's own claim rather than to a Chapter 3 Link's.

The one structural rule this chapter adds, and did not need for a design decision alone, concerns a continuation claim: whether a wave should carry on as it stands. Here, at least one challenger **MUST** stand outside the wave's own chain of accountability — the named independent reviewer designated at the wave's outset, or someone with equivalent independence. Everywhere else in this chapter, and everywhere in the series, a challenger's standing rests on subject accountability: whoever answers for the claim in practice is best placed to test it. A continuation claim breaks that pattern, because the documented failure it guards against is not incompetence but self-justification by the party responsible for the original choice, as §2's escalation and de-escalation findings state in full.

The chapter states this as a property of the challenge record — who put the challenge, and on what standing — never as a procedure for convening a meeting. A reader who was not in the room can then tell a continuation claim tested by an independent reviewer apart from one tested only by the wave that wants it to continue.

7.3 The nine challenges

The nine challenges now sit at §5.12, immediately after Benefits realisation, the last Move they test against.

7.4 Failure modes

Five failure modes recur across the drills this method was built from, carried here under their existing names and mapped to the Move each guards against, restated generically: no company, no figure, no artefact from any course case.

- **d-BANG — the big-bang cutover.** A single full-scope switch with no stated reversibility condition and no named authority who could abort it, aimed at the part of the business the programme exists to protect. E4 guards against it: a rollout posture naming neither a reversibility condition nor an abort authority is a schedule with a date attached, not a posture, and is not recorded.
- **d-THEATER — the register kept for a slide, not a decision.** An artefact that looks like a live risk register — rows, owners, dates — but carries no trigger written before the fact and is never reopened between reviews. E6 guards its creation; the weekly exposure review guards its use. What separates it from a working register is never appearance; it is whether the register is opened and updated between reviews.
- **d-SPONSOR — the checkbox sponsor.** A named executive who appears once, at launch, and is recorded as present rather than as the holder of a decision only they can make. E8 guards against it: each entry in the sponsor commitment record names one person and the one decision or resource commitment only they can authorise. A role, a committee or

a signature does not satisfy the entry.

- **d-COMMS — change mistaken for communication.** A plan built from messages, sessions and a channel — not wrong, and nowhere near sufficient, because none of it would ever read differently if nobody changed what they did. E7 guards against it by naming the dimension of resistance or readiness the messages do not reach; E9 guards against it by measuring the behaviour the messages do not change, refusing attendance or survey completion as evidence of it.
- **d-BENEFITS — "it shipped, so it is done."** A closeout complete on delivered outputs and silent on any re-measured baseline, often compounded by counting a pilot's numbers as the rollout's own. E11 guards the ledger that would show the gap; E10 guards GC3, which tests the claim before another wave is funded. This is the chapter's climactic failure mode: every other one can survive a single weak Move, and this one survives the whole chain if the ledger and GC3 are not both checked.

7.5 What this paper does not claim

No defensible failure or success base rate for change programmes exists in either direction. The widely repeated change-failure narrative has no valid empirical support (Hughes, 2011), and no study located since has supplied a rival figure for either failure or success. This chapter prints no number, in either direction, and for exactly that reason tests each programme's own claims one at a time — a sponsor named, a band carrying an action, a benefit carrying a baseline — rather than arguing from a population statistic that does not exist to argue from.

Nothing here asserts that a programme executed this way succeeds more often, decides better, or realises more benefit than one executed otherwise. What this chain buys is auditability — that a sponsor was a person, a band carried an action, a benefit carried a baseline, a continuation claim cited a number — each checkable by someone who was not in the room. Traceability buys auditability, never decision quality.

The practitioner canon this chapter draws its vocabulary and failure-mode names from is discussed in full at §6.2. It supplies naming, never the authority of a rule; no MUST anywhere in this chapter rests on any of it. Where one cross-country empirical study exists on a neighbouring question — benefits-management practice and programme success — it found associations with value-creation measures specifically, weak and inconsistent overall (Serra & Kunc, 2015), and is cited on that narrower finding, not a broader one.

An audited public case named here as a failure-mode illustration is one example, never a statistical sample of programme performance. Where this chapter's strongest structural claim rests on an analogy — mirroring tested on codebases, reference-class thresholds assuming a comparable population, escalation research drawn from decisions generally, not transformation governance — it is stated as an extension, not a validated transplant. Separating kill authority from a wave's own chain of accountability is a reasoned bet on evidence that de-escalation is usually triggered from outside, not evidence that this design produces more kills. A programme that has withstood every challenge in §7.1 is challengeable, not thereby well run.

Most of this register's requirements are MUST-level, and not one of them is evidence that a programme run this way succeeds more often. What each buys is a specific, named thing a later reader can check — that a sponsor was a person, that a band carried an action, that a benefit carried a baseline. The register stays silent on the parts of execution nobody has measured: whether these checks change what anyone decides, whether an independent reviewer actually kills a wave the sponsor wants, and whether a traceable ledger survives a year of reorganisation. Naming those residual exposures is part of the method, not an apology for it.

This paper is not advice, and it is free to use.

8. Applicability

Eleven Moves, a three-verdict entry contract and a three-tier governance cadence cost more than a status meeting and a stakeholder list. That cost buys something specific — a wave needing only a technology rollout, with no funding decision and no behavioural change at stake, does not owe it.

Use it when

- The released pack already carries all five HX1 components, and more than one wave's funding is still to be decided. GC3 tests continuation against tracked movement; an already-funded piece of work gives it nothing left to test.

- More than one team or supplier is implicated in how the work will flow. A sourcing decision and a delivery-boundary record answer different questions, and only a forced check catches a boundary drawn around a supplier contract rather than the change itself.
- Behavioural change is required in more than one stakeholder group, not confined to operating a new tool. A stakeholder read, a sponsor commitment and an adoption measure each refuse a different failure a merged "people" step would let through.
- The organisation is willing to name, before spend is authorised, a sponsor, a reviewer outside the wave's own chain of accountability, and the clearing metric each transition will be tested against.
- The programme runs long enough, and exposes enough people, that a standing weekly and monthly review earns its own mandate, ahead of GC3's quarterly review.

Lighten it or choose another approach when

- The change is confined to a single team and supplier, with no stakeholder group's behaviour to change — no funding decision, no sourcing decision, no adoption measure worth naming. Running the full chain then costs more than the change is worth.
- The pack this chapter would receive does not yet carry all five components Chapter 3's exit contract names, or the boundary marker does not say what this chapter owns. A RETURNED verdict at the entry contract is more defensible than running Moves against an incomplete pack.
- The organisation cannot or will not name a sponsor, a mandate and a funding envelope separately from who executes the work. Running the later Moves without that separation produces records with nothing to check.
- No wave transition yet has a named clearing metric, value and date. The governance cadence has nothing to test without one, and a review against an unset threshold states nothing falsifiable.

Preconditions. Before the entry contract runs, confirm:

1. an HX1 pack carrying all five components Chapter 3's exit contract names, with the HX2 boundary naming what this chapter owns, or a clear route to close that gap first;
2. standing to name an accountable sponsor, a recorded mandate, a funding envelope and a reviewer outside the wave's chain of accountability — not a single approver who also delivers the work;
3. willingness to state, before the first wave's funding is authorised, the clearing metric, threshold and date GC3 will test. This chapter mints no metric family of its own, tracking Chapter 3's own tree in flight instead;
4. a route back to Chapter 3's published referral that is genuinely available, not merely notional, for a pillar whose assumption is falsified once work is under way.

Wanting a wave funded this quarter and clearing every precondition first are not always compatible. The transparent response is to scope down deliberately — fewer Moves run in full, more of the pack referred back across HX2 — rather than clearing a precondition quietly and calling the wave released.

9. Why this decomposition and not the alternatives

Eleven Moves are not the only split the material admits, and the case for eleven rests on what each split lets a single card refuse — not on comparing external methods this research did not re-open, and not on a measured result no source has produced for any ordering.

Against keeping the training material's own ten-unit split. Its fifth unit carries two artefacts under one code — a generation technique and a record — with different admission conditions (one needs a plan not yet committed; the other needs candidate causes) and different evidence bases: a laboratory finding about the number of reasons a person generates, set against a mathematical critique of ordinal scoring matrices. One card cannot carry two contract checks without hiding one of them. The split into a separate Pre-mortem Move and Risk register Move is forced by the material, not chosen for symmetry, and it is why this chapter carries eleven operations, not ten.

Against a single change-management Move. Stakeholder diagnosis, Sponsorship and coalition, and Adoption would collapse into one "people" step in most practitioner sequences. They resist collapse because their refusals differ: one refuses an unnamed dimension of a stakeholder read, the next a named role standing in for a named person, the third exposure recorded in place of measured behaviour. Merged, a programme could satisfy the easiest refusal and never meet the other two — precisely the communication-only failure this chapter names.

Against folding the governance cadence into benefits realisation. A funding-continuation decision and a realisation measurement look adjacent and are, on inspection, opposed. The first is taken mid-flight, by people invested in the wave continuing; the second is taken after the fact, against a baseline nobody holding that stake may move. Merging the two would place the party that wants a favourable outcome in charge of the record that could refuse one.

Against merging sourcing into delivery boundaries. A sourcing decision and a boundary record answer different questions — who builds the work, and where it is allowed to flow — and each cites a different parent artefact. Merged, a boundary could be justified by a supplier contract rather than the shape of the change itself, reintroducing by another route the structure-change mismatch this chain exists to catch.

Against taking a single practitioner change model as the spine. An eight-step model built from retrospective case narrative and a five-stage model published by the vendor that sells it are both usable as vocabulary and neither is testable as a sequence: a spine whose own steps cannot be failed is not a chain. This chapter takes the vocabulary such models offer and builds its sequence only from what a contract check can refuse.

Against a longer chain. Every further split considered — scheduling apart from mobilisation, dependency management apart from the risk register — produced a further Move with no refusal distinct from its neighbour's. A Move that cannot refuse anything is a heading, not an operation.

None of the six arguments claims that eleven Moves execute more reliably than ten, or any other count; each states only what a given card's contract check would otherwise be unable to refuse. This decomposition is reasoned, not empirically benchmarked, and no ordering — this one included — has been measured against another on outcomes.

Extended rationale: foundations and the contribution's own bounds

Staw's (1976) laboratory study found that participants who had personally chosen a course of action later committed more resource to it, once it began failing, than participants who had not made the original choice — the escalation-of-commitment finding. Brockner's (1992) review extends the mechanism: self-justification, sunk-cost framing and a threat to the decision-maker's own competence combine to make a chosen course of action harder to abandon than an inherited one. Keil and Robey's (1999) study of troubled technology projects found the original decision-maker structurally unreliable as the sole source of its own reversal.

Flyvbjerg's (2008) reference-class forecasting corrects a plan's own inside-view estimate against comparable past projects rather than trusting it on its own case; Flyvbjerg, Garbuio and Lovallo (2009) frame cost and benefit forecasts as distorted by delusion and deception alike, trusted only once tested against outside evidence. Cantarelli, Davis, Pinto and Turner's (2025) counterweight treats the choice of reference class as a judgement call, so a corrected estimate informs a decision a named person still makes, never a threshold.

Conway's (1968) proposition that structure mirrors the organisation that produced it is supported, for codebases, by MacCormack, Baldwin and Rusnak's (2012) comparative study; this chapter extends the same shape to a wave's own boundaries, where a boundary cut along the wrong seam reproduces a mismatch this literature already names.

What this chapter's own principle adds, beyond these borrowed foundations, is the arrangement, not the parts: a three-verdict release contract with two repair addresses, operational review kept separate from funding authority, a reviewer named outside the wave's own chain of accountability, a register tying a likelihood-and-severity read to a pre-declared rule, and a benefits ledger that keeps planned apart from realised.

10. Scope, roadmap and references

10.1 Scope of Chapter 4

Chapter 4 specifies the Execution Chain in full: the eleven Moves, E1 through E11, across the four acts; the entry contract at the HX2 boundary and its three verdicts; the two-sided contract-check pattern, including the one admission check E10 adds; the GC1–GC3 cadence and the boundary that reserves funding release to GC3 alone; the banded decision-rule register; the nine challenges and the five named failure modes; and the exit contract that closes on a benefits ledger checked against Chapter 1's own success criteria. It fixes no total across the rules it states, and it mints no metric family of its own — where it tracks Chapter 1's tree in flight, an execution measure that tree does not already carry is recorded and reported as an execution measure, never folded into it.

A companion practitioner guide — worksheets, detailed procedures and a full worked case — is in preparation. It is scoped separately, is not part of this specification, and no rule stated here waits on it or is incomplete without it.

Three things this chapter does not carry, each by a decision already recorded rather than an omission:

- **No named case.** Every illustration in this chapter is generic, unnamed and carries no financial figure; a worked instance with people, dates and outcomes belongs to the practitioner guide, not to a specification.
- **No Move-to-Move parent register.** Chapter 3 publishes a full worked derivation register as its own Appendix B; a Chapter 4 counterpart would need the instance values the previous bullet already rules out, so this chapter carries §5.13 alone.
- **No Russian edition.** Translation is queued behind Chapter 3's own, which itself awaits acceptance of that chapter's English text.

10.2 Roadmap

The practitioner guide named in §10.1 follows this specification on its own schedule; nothing here depends on its arrival, and nothing in it may revise a rule this chapter states. A Russian edition follows Chapter 3's. All nine challenges ship in this chapter; nine is fixed as this chapter's ceiling — a deliberate closure of the set, not a gap towards a tenth added later. A future revision might extend the banded decision-rule register or draw the structural, value-free Move-to-Move register the previous section declines to carry now. None of this is a commitment, and this chapter states no date against any of it.

10.3 Changelog

Date	Change
22 September 2026	First published — the specification of the Execution Chain, its eleven Moves and four acts.
28 September 2026	§7.3's nine challenges moved to §5.12, immediately after §5.11 Benefits realisation, verbatim, numbering 1–9 unchanged; Figure 7.1 renumbered Figure 5.12 (caption, SVG <title> and alt text) to match; Appendix A's 37-row table moved to new §5.13, immediately after §5.12, whole and verbatim, and the now-empty "Appendix A" heading removed; §2 Foundations and §3.4 condensed, with the fuller escalation, reference-class and mirroring account moved into a new "Extended rationale" sub-heading inside §9; §1.2's RFC 2119/8174 boilerplate paragraph replaced with a pointer to the shared conventions page; repeated citation recaps at §7.2, Challenge 1 and §7.5 condensed to cross-references, each pointing to the one place its finding is now stated in full; §3.2's Chapter 3 citation corrected from "§6.5" to "§5.11". §7.1's opening sentence corrected to point to the challenge set's new location (§5.12) instead of announcing it as following; Figure 7.1's legend corrected to match its updated caption. No Move, admission check, contract check, challenge or verdict was added, removed or reworded, and every rule sentence keeps its original wording.
28 September 2026	Version numbers removed; changes are now recorded by date; no rule changed.

This changelog belongs to this chapter alone. It advances on its own schedule, independently of Chapters 1 through 3, and an entry here implies nothing about any sibling chapter's own release. A future change to this chapter is recorded as a new entry here, never as a silent edit to this one.

10.4 References

Listed alphabetically; cited in-text by author and year, matching Chapters 1 through 3's convention.

- S. H. Appelbaum, S. Habashy, J.-L. Malo and H. Shafiq, "Back to the Future: Revisiting Kotter's 1996 Change Model" (*Journal of Management Development* 31(8), 2012) — <https://doi.org/10.1108/02621711211253231>
- J. B. Barney, "Firm Resources and Sustained Competitive Advantage" (*Journal of Management* 17(1), 1991) — <https://doi.org/10.1177/014920639101700108>
- S. Bradner, "Key Words for Use in RFCs to Indicate Requirement Levels" (BCP 14, RFC 2119, RFC Editor, 1997) — <https://www.rfc-editor.org/rfc/rfc2119>
- J. Brockner, "The Escalation of Commitment to a Failing Course of Action: Toward Theoretical Progress" (*Academy of Management Review* 17(1), 1992) — <https://doi.org/10.5465/amr.1992.4279568>

- C. C. Cantarelli, K. Davis, J. K. Pinto and N. Turner, "Reference Class Forecasting: Promises, Problems, and a Research Agenda Moving Forward" (*Production Planning & Control*, 2025) — <https://doi.org/10.1080/09537287.2025.2578708>
- M. E. Conway, "How Do Committees Invent?" (*Datamation* 14(5), 1968) — https://www.melconway.com/Home/Committees_Paper.html
- L. A. Cox, "What's Wrong with Risk Matrices?" (*Risk Analysis* 28(2), 2008) — <https://doi.org/10.1111/j.1539-6924.2008.01030.x>
- S. Cummings, T. Bridgman and K. G. Brown, "Unfreezing Change as Three Steps: Rethinking Kurt Lewin's Legacy for Change Management" (*Human Relations* 69(1), 2016) — <https://doi.org/10.1177/0018726715577707>
- A. C. Edmondson, "Psychological Safety and Learning Behavior in Work Teams" (*Administrative Science Quarterly* 44(2), 1999) — <https://doi.org/10.2307/2666999>
- B. Flyvbjerg, "Curbing Optimism Bias and Strategic Misrepresentation in Planning: Reference Class Forecasting in Practice" (*European Planning Studies* 16(1), 2008) — <https://doi.org/10.1080/09654310701747936>
- B. Flyvbjerg, M. Garbuio and D. Lovallo, "Delusion and Deception in Large Infrastructure Projects: Two Models for Explaining and Preventing Executive Disaster" (*California Management Review* 51(2), 2009) — <https://doi.org/10.2307/41166485>
- Google, *The Site Reliability Workbook*, "Canarying Releases" — <https://sre.google/workbook/canarying-releases/>
- R. Hirschheim and M. Newman, "Information Systems and User Resistance: Theory and Practice" (*The Computer Journal* 31(5), 1988) — <https://doi.org/10.1093/comjnl/31.5.398>
- D. T. Holt, A. A. Armenakis, H. S. Feild and S. G. Harris, "Readiness for Organizational Change: The Systematic Development of a Scale" (*Journal of Applied Behavioral Science* 43(2), 2007) — <https://doi.org/10.1177/0021886306295295>
- M. Hughes, "Do 70 Per Cent of All Organizational Change Initiatives Really Fail?" (*Journal of Change Management* 11(4), 2011) — <https://doi.org/10.1080/14697017.2011.630506>
- M. Keil and D. Robey, "Turning Around Troubled Software Projects: An Exploratory Study of the Deescalation of Commitment to Failing Courses of Action" (*Journal of Management Information Systems* 15(4), 1999) — <https://doi.org/10.1080/07421222.1999.11518222>
- W. C. Kim and R. Mauborgne, "Fair Process: Managing in the Knowledge Economy" (*Harvard Business Review*, January 2003) — <https://hbr.org/2003/01/fair-process-managing-in-the-knowledge-economy>
- G. Klein, "Performing a Project Premortem" (*Harvard Business Review*, September 2007) — <https://hbr.org/2007/09/performing-a-project-premortem>
- B. Leiba, "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words" (BCP 14, RFC 8174, RFC Editor, 2017) — <https://www.rfc-editor.org/rfc/rfc8174>
- D. Lovallo, M. Cristofaro and B. Flyvbjerg, "Governing Large Projects: A Three-Stage Process to Get It Right" (*Academy of Management Perspectives* 37(2), 2023) — <https://doi.org/10.5465/amp.2021.0129>
- A. MacCormack, C. Baldwin and J. Rusnak, "Exploring the Duality Between Product and Organizational Architectures: A Test of the 'Mirroring' Hypothesis" (*Research Policy* 41(8), 2012) — <https://doi.org/10.1016/j.respol.2012.04.011>
- M. L. Markus, "Power, Politics, and MIS Implementation" (*Communications of the ACM* 26(6), 1983) — <https://doi.org/10.1145/358141.358148>
- D. J. Mitchell, J. E. Russo and N. Pennington, "Back to the Future: Temporal Perspective in the Explanation of Events" (*Journal of Behavioral Decision Making* 2(1), 1989) — <https://doi.org/10.1002/bdm.3960020103>
- National Audit Office, *Over-Optimism in Government Projects* (2013) — <https://www.nao.org.uk/insights/optimism-bias-paper/>
- National Audit Office, *Review of the Final Benefits Statement for Programmes Previously Managed under the National Programme for IT in the NHS* (2013) — <https://www.nao.org.uk/reports/review-of-the-final-benefits-statement-for-programmes-previously-managed-under-the-national-programme-for-it-in-the-nhs/>
- S. Oreg et al., "Dispositional Resistance to Change: Measurement Equivalence and the Link to Personal Values Across 17 Nations" (*Journal of Applied Psychology* 93(4), 2008) — <https://doi.org/10.1037/0021-9010.93.4.935>
- C. K. Prahalad and G. Hamel, "The Core Competence of the Corporation" (*Harvard Business Review*, May–June 1990) — <https://hbr.org/1990/05/the-core-competence-of-the-corporation>
- C. E. M. Serra and M. Kunc, "Benefits Realisation Management and Its Influence on Project Success and on the Execution of Business Strategies" (*International Journal of Project Management* 33(1), 2015) — <https://doi.org/10.1016/j.ijproman.2014.03.011>

- M. Skelton and M. Pais, *Team Topologies* (IT Revolution, 2019)
 - B. M. Staw, "Knee-Deep in the Big Muddy: A Study of Escalating Commitment to a Chosen Course of Action" (*Organizational Behavior and Human Performance* 16(1), 1976) — [https://doi.org/10.1016/0030-5073\(76\)90005-2](https://doi.org/10.1016/0030-5073(76)90005-2)
 - S. Tadelis and O. E. Williamson, "Transaction Cost Economics", in R. Gibbons and J. Roberts (eds.), *The Handbook of Organizational Economics* (Princeton University Press, 2012) — <https://doi.org/10.1515/9781400845354-006>
 - B. C. Wintle et al., "Verbal Probabilities: Very Likely to Be Somewhat More Confusing than Numbers" (*PLOS ONE* 14(4), 2019) — <https://doi.org/10.1371/journal.pone.0213522>
 - O. Zwikael and J. Smyrk, "A General Framework for Gauging the Performance of Initiatives to Enhance Organizational Value" (*British Journal of Management* 23(S1), 2012) — <https://doi.org/10.1111/j.1467-8551.2012.00823.x>
-

About the author

Aliaksei "Alex" Ivanouski is a Fractional Chief Technology Officer (CTO) and founding engineer with 10+ years building fintech, adtech and backend products. A military-engineer-turned-founder, he holds a Master of Business Administration (MBA) with Distinction from Bayes Business School, City St George's, University of London, and completed the CTO Programme at the National University of Singapore. He built the Business Strategy as Code method and applies it in a practitioner programme at strategyascode.org.

Colophon

- **Status.** Published 22 September 2026; last updated 28 September 2026.
- **Production.** Generated from the author's source text.
- **How to cite.** Ivanouski, A. (2026). *Business Strategy as Code: An Auditable Method for Executing a Transformation Strategy*, Chapter 4. <https://whitepaper.strategyascode.org/execution-chain/>. Chapters 1, 2 and 3 are cited the same way, each with its own chapter number, title and URL.
- **Updates.** Changes to a chapter are recorded, dated, in its own §10.3; earlier wording is not silently altered, and one chapter's changelog never affects another's.