

# Business Strategy as Code: A Falsifiable Method for Strategic Diagnosis

Chapter 1 — Diagnosis

September 2026

Aliaksei Ivanouski

---

## Abstract

A team writes a crisp, sourced problem statement, an analysis chain traces every pillar back to it without one broken link, and the recommendation still answers the wrong question, because nothing put the statement itself under test. Business Strategy as Code (BSaC\_) — Diagnosis is a method for producing and testing that statement before an analysis chain may open on it, run as four transforms — D1 Frame, D2 Evidence, D3 Compose, D4 Test — each admitting its input under its own check and releasing its output only once that output meets its own contract check. A statement is admitted only once it has survived a fixed set of adversarial challenges, put by named human challengers who are competent by their own accountability for what the statement measures or costs — a named, obtainable observation that would show it false must exist, and must have been sought, before anyone commits to it. This paper specifies Chapter 1, the diagnosis stage, and marks each rule by level. It makes a diagnosis testable; no study shows that a better diagnosis produces a better outcome.

---

## About this document

- **Language.** English.
- **Illustrations.** Generic and informative: no company names, and no fixed count beyond the method's own — four transforms, four verdicts, nine challenges.
- **Nature.** A method description for consultancy and advisory work, not legal, financial or investment advice.
- **Status.** Published 22 September 2026; last updated 28 September 2026.
- **Companion guide.** A practitioner guide with detailed procedures, worksheets, question scripts and a full worked case is in preparation.
- **Relationship to Chapter 2.** This chapter specifies how a diagnosis produces the single statement that Chapter 2's analysis chain admits at its entry contract (Chapter 2, §3.2); §3.2 below states the principle that governs how the two chapters connect. Reading this chapter does not require Chapter 2, though the chain has nothing to open on until a diagnosis has produced one.

---

## 1. Introduction

### 1.1 The problem and the use case

An analysis can trace every pillar back through its coded stops to a named source and still be a careful analysis of the wrong problem: the entry contract admits a statement on its shape, not on its truth. Chapter 2 fixes what an admitted statement must contain — a named decision-maker, explicit scope, a sourced baseline, minimal context and numeric success criteria, with every clause carrying its own evidence (Chapter 2, §3.2) — and a statement can satisfy every one of those fields, fluent, sourced, signed by a willing decision-maker, and still describe a symptom rather than the difficulty producing it. A field that is missing or unsourced is a shape failure the entry contract catches directly; a field that is present, sourced and still wrong is not caught there at all. Producing a statement that passes that further test, and recording how it was tested, is what this chapter specifies.

This chapter serves whoever is asked to state a problem before an analysis of it begins — consultants and advisers scoping an engagement, strategy and transformation leads opening a piece of work, and the decision-makers and delegates who will be asked to commit to what the statement claims. It fits the same decisions Chapter 2 fits — costly to reverse, facing outside scrutiny, or changing hands. It does not replace judgement in naming a problem; it records how the naming was tested, and by whom, so that someone other than its author can check it.

## 1.2 Conventions

Requirement-level key words (MUST, SHOULD, MAY) follow the shared conventions at [whitepaper.strategyascode.org/conventions/](http://whitepaper.strategyascode.org/conventions/). Here they are fixed independently of Chapter 2, and a MUST binds the diagnosis.

Chapter 2 fixes "gate" and "stop" as its own vocabulary (eight quality gates, fifteen numbered units); this chapter reuses neither word, calling its own adjudications challenges and its own four units transforms.

This chapter's four transforms are written bare — D1, D2, D3, D4 — never with a middle dot, which belongs to Chapter 2's coded families (its internal-fact family, I-1 ... I-n, with different letters — i, j, k, m — naming different codes, n never a count). This chapter mints no code family of its own and uses that notation only when naming Chapter 2's family by Chapter 2's own name.

A **transform** is one of the diagnosis's four operations — D1 Frame, D2 Evidence, D3 Compose, D4 Test — each of which MUST admit its input under its own admission check and emit its output under its own contract check. An **admission check** is the condition an input MUST meet before its transform will act on it; where the condition is not met, the transform does not run, and returns a named deficiency instead. A **contract check** is the condition a transform's output MUST meet before it may be cited downstream — used in exactly the sense Chapter 2 already publishes, not redefined here. A **challenge** is one of the nine fixed, adversarial tests D4 runs against a composed statement, each with its own pass condition and level. A **challenger** is a person with recorded standing who puts a challenge to the statement, neither its composer nor its named decision-maker. **Standing** is the recorded reason a specific challenger is competent to attack this specific claim — accountability for the metric or the consequence the statement asserts, never seniority alone. The **challenge record** is the diagnosis's audit artefact: one row per challenge naming the question, the challenger, the answer, the evidence cited and the outcome, retained with the statement together with any superseded framing. A **verdict** is the single, fixed-vocabulary outcome D4 issues on a diagnosis — ADMITTED, PROVISIONAL, DEFICIENT or REFRAME — decided on the diagnosis as a whole, never a single evidence item. **Ratification** is the act of a statement being allowed to stand, the opposite incentive to a challenge, never performed by the challenger who ran a challenge against the same statement. Chapter 2 does not use the word, and this chapter requires no ratification artefact, because the ratification fact — who committed to the target, and by when — is already carried at MUST by challenge 3. **The frame** is D1's artefact: a named decision-maker, scope stated in and out, the situation, the complication and the consequence — a stable, solution-free and cause-free account of what is happening, held apart from the evidence that will test it. A **delegate** is a person, other than the named decision-maker, who may propose a numeric target on the decision-maker's behalf when the decision-maker cannot yet commit; a delegate's proposal caps the resulting verdict at PROVISIONAL until the named decision-maker commits in their own right (§5.3). **The statement** is D3's artefact: the single candidate sentence, with its clause-to-evidence map and numeric success criteria, that becomes the problem statement Chapter 2's entry contract admits once ADMITTED; it is the only artefact this chapter hands to Chapter 2 (§3.2).

## 1.3 How to read this paper

- §2–§4 give this chapter's foundations, principle and four-transform shape; §4.3's table gives every transform's key rules at a glance.
- §5 gives each transform as a card — admission check, correct use, contract check — plus §5.5, the nine challenges.
- §6 covers evidence and the handoff to Chapter 2; §8 applicability; §9 the alternatives considered and the extended rationale; §10 scope, roadmap and references.

## 1.4 Abbreviations

This chapter mints no codes of its own (§6.2); the short forms below are the only abbreviations its prose uses.

- BCP — Best Current Practice
- BSaC\_ — Business Strategy as Code
- CTO — Chief Technology Officer
- KPI — Key Performance Indicator
- MBA — Master of Business Administration
- RCA — root-cause analysis
- RFC — Request for Comments
- SCQA — Situation, Complication, Question, Answer

## 2. Foundations

Diagnosis draws on four bodies of practice, distinct from the four Chapter 2 draws on for the analysis chain — requirements traceability, architecture decision records, evidence-based management and Rumelt's kernel (Chapter 2, §2).

Lyles and Mitroff (1980) treat problem formulation as a studiable process in its own right; diagnosis borrows that premise, not a finding about outcomes. Rumelt (2011b) opens the kernel of good strategy with diagnosis, naming the challenge before any guiding policy; Wedell-Wedellsborg (2017) found most surveyed executives move into solution mode before checking their own understanding. Nickerson (1998), on confirmation bias, is the reason this chapter puts a frame to challengers who did not write it. Locke and Latham (2002), and the Ordóñez, Schweitzer, Galinsky and Bazerman (2009)/Locke and Latham (2009) dispute over prescribing goals, together support only a narrower point: a target that is not numeric, dated and owned cannot later be shown to have been met or missed.

None of these sources examines a diagnosis run this way; each is cited for the narrower point it establishes. §9's extended rationale sets out the fuller argument.

## 3. The principle

### 3.1 The rule, one stage earlier

The chain rests on one rule: no output without a traceable input, tested by whether every citation path ends at a root (the analysis chain, Chapter 2, §3.1). Diagnosis sits one stage earlier than any stop the chain runs, and the rule holds there too, before its usual instrument exists — nothing here mints a code yet, so what travels instead is provenance: a named decision-maker's own account, a metric traced to a named system, a market claim traced to a named report.

The requirement does not relax for the absence of codes; it changes currency. What this stage produces, once every check has passed, is a single candidate sentence — its clause-to-evidence map, its numeric success criteria — referred to throughout this chapter as the statement, the only thing this stage hands forward.

### 3.2 The exit contract

"Stop 1 MUST NOT begin without a scoped problem statement." — the analysis chain, §3.2

The analysis chain (Chapter 2, §3.2) publishes what the statement must contain: five prerequisites, sourced evidence for every clause, six quality criteria and the single-sentence form. This chapter does not restate those fields; it specifies how one is produced.

Where this chapter and the entry contract speak to the same requirement, the entry contract governs. This chapter MUST NOT weaken a §3.2 requirement and MUST NOT restate its fields; where it adds a requirement §3.2 does not make, it says so. Every addition — a challenger drawn from outside the statement's own composer, a delegate marked as such, a scope stated across facets §3.2 leaves general — is strictly stricter than the interface it sits beside, never looser.

Every citation path terminates at a root backed by a named external source, a registered internal fact, or the entry problem statement. As §6.2 states, diagnosis mints no code family of its own; its trail terminates at the statement, which remains a root of the chain's graph, never a node derived inside it. The clause-to-evidence map is this chapter's own retained artefact, not part of the chain's derivation graph.

**The problem statement is the only connection between the chapters.** There is no side-channel artefact and no second interface. An organisation's own internal facts gathered while composing the statement MAY be registered under Chapter 2's own internal-fact codes, I-1 ... I-n, at Chapter 2's own MAY level. §6.3 gives the mechanism in full, and the reason it is not a second interface. An I-code belongs to the chain's own family, registered under the chain's own MAY-level rule, and records a fact about the organisation, not a fact about this chapter's method. The statement remains the only connection this method defines.

### 3.3 Two-sided contracts

Chapter 2's stops need only a contract check, since each takes the already-checked output of the stop before it. Diagnosis has no such predecessor: D1 takes a decision-maker's stated concern, not a checked artefact, with nothing upstream to have vetted it. That is the failure this chapter's one structural addition exists to close: every transform

MUST declare not only a contract check on what it emits, but an admission check on what it takes. An input that fails admission is refused, named and returned — never transformed into a plausible-looking output.

The two sides recur asymmetrically through all four transforms: a concern admits at D1 against a nameable decision-maker and something that has changed; a classified evidence-demand list admits at D2; an evidence set with no gap against a required field admits at D3; challengers with recorded standing admit at D4.

### 3.4 What is and is not the contribution

None of this chapter's material is new; a five-field intake, a named-owner check and a rule against smuggling a cause into a symptom's description already circulate as practical craft, and this chapter claims none of it as its own invention. What it contributes is the discipline that turns that craft into something able to fail loudly, set out in full at §9's extended rationale. The claim stops there: nothing here asserts that a diagnosis built this way reaches a better decision than one built any other way, only that it can be shown to have survived a genuine attempt to break it.

## 4. The diagnosis

### 4.1 Four transforms

Diagnosis runs as four transforms, each admitting its own input and emitting its own output under its own two checks: D1 Frame, D2 Evidence, D3 Compose, D4 Test.

D1 turns a decision-maker's stated concern into the frame — a stable, solution-free and cause-free account of what is happening, to whom and since when — together with an evidence-demand list naming what the frame will need. D2 turns that list into a sourced evidence set, together with a recorded search for whatever would disconfirm the frame rather than support it. D3 turns the evidence set into the statement itself: one sentence, its clause-to-evidence map, and numeric success criteria stated as a baseline, a target and a date. D4 puts that statement to nine fixed challenges and closes the diagnosis with exactly one verdict.

Each transform's output becomes the next transform's raw material; only D4's output crosses the diagnosis stage's own boundary — a single verdict, attached to a single statement.




**Figure 4.1 — The four transforms and their two-sided contracts.** Each transform checks what it admits as well as what it emits, and only one exit reaches the chain. This method's format.

### 4.2 Why four, and why this order

Four is not an arbitrary split, and the order is not incidental.

A single flat contract — one list of fields a statement must satisfy — cannot refuse an input before transforming it, cannot locate which of the four things behind it (the frame, the evidence, the composition, the challenge) is where a deficiency actually lives, and cannot distinguish "the evidence is missing" from "the statement is wrong" — different findings calling for different remedies: the first sends the diagnosis back to D2, the second forces a reframe.

The transforms SHOULD run in this order — D1, then D2, then D3, then D4 — because each needs what the one before it already produced. D2 cannot bound its own search without D1's evidence-demand list. D3 cannot build a clause-to-evidence map before an evidence set exists to map against. D4 cannot put an unfinished statement to challenge.

Two mergers were considered and rejected. Composing and testing stay separate because whoever composed a statement is the worst-placed person to find what would break it. Framing and evidence-gathering stay separate because D1's demand list is what makes D2 a finite task rather than an open-ended one.

### 4.3 The diagnosis at a glance

All four transforms are this method's own; §5 gives each in full, its admission check, its transform and its contract check stated as a paragraph rather than compressed into a table cell. The table below gives the shape of all four side by side, with each transform's key rules stated by content rather than by an internal identifier.

| Transform   | Admission check                                                                                                                   | Takes                                      | Produces                                            | Contract check                                                                                                                                  | Key rules (level)                                                                                                                                  |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| D1 Frame    | A live decision; a nameable role; a stated situation; something that changed                                                      | Concern; trigger; access to internal facts | Frame; evidence-demand list                         | Decision-maker named; scope stated and sufficient; solution-free and cause-free; consequence stated; demand list classified                     | Decision-maker named (MUST). Scope sufficient (MUST). No solution, no cause (MUST). Consequence stated (MUST). Demand-list item classified (MUST). |
| D2 Evidence | A classified evidence-demand list                                                                                                 | Frame; demand list                         | Evidence set; disconfirming search; gap list        | Baseline sourced; context sufficient; every item sourced; gaps named against §3.2 fields                                                        | Baseline sourced or gapped (MUST). Every item sourced (MUST). Disconfirming search recorded (SHOULD). Gap named against its field (MUST).          |
| D3 Compose  | No gap against a §3.2 field; a target committed, or proposed by a named delegate (proposal caps the verdict at PROVISIONAL, §5.3) | Evidence set; frame                        | Statement; clause-to-evidence map; success criteria | Every clause resolves; no overreach; criteria numeric with baseline, target and date; scope recoverable; six criteria adjudicated               | Every clause resolves (MUST). No overreach (MUST). Numeric criteria with a date (MUST). Target committed or marked proposed (MUST).                |
| D4 Test     | Challengers with recorded standing, none of them the composer or the decision-maker                                               | Statement; map; criteria                   | Challenge record; verdict                           | All nine challenges recorded with an outcome and a named challenger; no unanswered recorded as withstood; one verdict, with its required naming | All nine recorded (MUST). Unanswered never withstood (MUST). One verdict (MUST). ADMITTED requires a human challenger of record (MUST).            |

Read left to right, a row states what a transform needs already to be true before it will act, and what it must be able to show once it has acted. Read top to bottom, the "produces" column of one row is the raw material the "takes" column of the next names — the sequence §4.2 defends, laid flat. A fifth relationship sits beside the four transforms, not among them: the internal-fact hook (§6.3).

### 4.4 The four verdicts

D4 closes each diagnosis with exactly one verdict, drawn from a fixed vocabulary of four words, never a fifth invented for the occasion.

ADMITTED requires that the statement meet §3.2, withstand all nine challenges before a human challenger of record, and carry a target the named decision-maker has committed to. The record names the challengers and their standing, and names who committed to the target and by when.

DEFICIENT names a required field the diagnosis could not evidence, or names that no challenger with standing could be found. Either way, the chain does not begin. REFRAME means the statement survived testing but was found to describe a symptom of a different problem. The record names the evidence item that forced the change and the framing now superseded, and a superseded framing is never re-adopted. The history of reframes stays acyclic — the same rule

the chain already publishes for its own citations, reused here rather than reinvented. Where no framing survives the challenge set at all, the verdict **MUST** be **DEFICIENT** once more, naming the field that still cannot be evidenced. The diagnosis **MUST NOT** loop past that point.


A diagnosis carries exactly one pending state, not two. **PROVISIONAL** means the statement is complete and has already survived challenge, but one required human act has not yet happened. The record names which act — a challenge a challenger with standing has not yet put, or a target no one has yet committed to — and that act **MUST** be completed before stop 1 begins. What distinguishes one **PROVISIONAL** diagnosis from another is always the named act still missing, never a second word invented for a second kind of pending.

None of the four words is a courtesy read of how the diagnosis went; only **ADMITTED** crosses into the chain.

## 5. The four transforms

All four transforms are this method's own. What follows gives each in full, as a card: its purpose, its admission check, its correct use, its contract check, its most common misuse and when to adapt or skip it.

### 5.1 D1 — Frame



**Attribution.** This method's own transform, not adapted from a single named framework. Its consequence test is the sibling of the consequence check Chapter 2 calls a Monday test (Chapter 2, gate 1), given its own name here because Chapter 1 mints no gate of its own.

**Purpose and position.** D1 turns a decision-maker's stated concern into a frame: a stable account of what is happening, to whom and since when, held apart from any proposed solution and any asserted cause. It is the transform every other transform depends on — D2's evidence-demand list, D3's clauses and D4's challengers all trace back to a field this transform names first.

| Field                 | Placeholder                 | Evidence demanded               |
|-----------------------|-----------------------------|---------------------------------|
| <b>Decision-maker</b> | <i>named role</i>           | role, not a person or committee |
| <b>Scope in</b>       | <i>included boundary</i>    | sufficient for downstream use   |
| <b>Scope out</b>      | <i>excluded boundary</i>    | sufficient for downstream use   |
| <b>Situation</b>      | <i>stable state, dated</i>  | no solution stated              |
| <b>Complication</b>   | <i>what changed, dated</i>  | no cause stated                 |
| <b>Consequence</b>    | <i>action if frame true</i> | stated and non-trivial          |

*Six fixed fields — not a repeating code family.*

**Figure 5.1 — The frame: what D1 must hold.** Each row carries the evidence it will demand of D2. This method's format.

**Correct use.**


- One decision-maker **MUST** be named by role; "the organisation", a committee or a vacant role is not one. Where no role can be named, D1 does not transform: it returns DEFICIENT, naming the decision-maker field. The check is on nameability, not availability — a decision-maker who cannot be reached today is still named; one who cannot be named at all stops the transform before it starts. The delegate case sits downstream, at D3, not here.
- Scope **MUST** be stated as in and out, and **MUST** be sufficient for the statement to be used downstream without any further artefact; it **SHOULD** cover four facets — organisational unit, offering or arena, geography, horizon.
- The frame **MUST** state what is happening, to whom and since when, and **MUST NOT** contain a proposed solution or assert a cause.
- The consequence test asks one question: what would the named decision-maker do differently if this were true? The frame **MUST** state an answer that is non-trivial.
- D1 **MUST** emit an evidence-demand list, every item classified as proprietor-supplied or externally observable.

**Most common misuse.** *Solution in disguise* — a proposed fix stated as if it were the situation, so the frame already contains its own answer.

**When to adapt or skip** (SHOULD NOT skip; MAY adapt). D1 **SHOULD NOT** be skipped, because every downstream demand traces to a field it names. Situation and complication **SHOULD** be stated separately, because a frame that fuses the two obscures which half changed. Where the decision-maker holds delegated authority — a mandate delegated to them from elsewhere, distinct from a delegate's target proposal at §1.2 and §5.3 — that authority is named in the scope field; the frame is not weakened by it.

**Illustration (informative).** An operations lead names falling on-time delivery as the concern. The frame states the drop, to which routes, and since when; it holds out any claim about which carrier caused it, and states that the lead would re-tender a route contract if the frame were true. An earlier draft naming "switch carriers on the affected routes" is rejected: it is a solution, not a situation.

## 5.2 D2 — Evidence



**Attribution.** This method's own transform, not adapted from a single named framework. Its disconfirming-search discipline draws on Nickerson's (1998) account of confirmation bias, cited here as rationale for why the search is asked for, not as this transform's origin.

**Purpose and position.** D1 states what must be known; D2 goes and gets it. It obtains every fact the frame demanded, with its source, searches for evidence that would disconfirm the frame rather than only what supports it, and records where a demanded fact cannot be found. Its evidence set is what D3 composes from and what D4 will later ask a challenger to trace.

| Item    | Class                                                  | Source              | Bearing                                      |
|---------|--------------------------------------------------------|---------------------|----------------------------------------------|
| item ·1 | <i>proprietor-supplied /<br/>externally observable</i> | <i>cited source</i> | <i>supports / challenges /<br/>/ neutral</i> |
| item ·2 |                                                        |                     |                                              |
| ⋮       | ⋮                                                      | ⋮                   | ⋮                                            |
| item ·n |                                                        |                     |                                              |

#### disconfirming search

*⋮ = the range continues; n is not a count.*

*The disconfirming search is recorded once, as its own row outside the ·1 ... ·n range.*

**Figure 5.2 — The evidence set: one row per item.** Class decides who can supply an item and whether its absence halts the transform. This method's format.

#### Correct use.


- Every evidence item **MUST** name its source; it **SHOULD** be specific and dated.
- Each item **SHOULD** be marked as supporting, challenging or neutral towards the frame, so a later reader can see the disconfirming search was not a formality.
- A baseline metric **MUST** exist with its source; where none does, the gap is recorded and instrumentation named as the first work.
- D2 **SHOULD** record a deliberate search for evidence that would disconfirm the frame, and its result, whether or not anything was found.
- Market and customer context **MUST** cover what the frame's clauses need, and **SHOULD** go no further.
- Two classes of fact meet different tests here. A proprietor-supplied fact — one only someone inside the organisation can produce — blocks the transform if it cannot be obtained: its absence is a gap against a named frame field, and D3 will refuse to admit past it. An externally observable fact **SHOULD NOT** halt the transform if it cannot be obtained; it is recorded instead as an open research item, because the world, unlike the organisation, owes the diagnosis no answer.

**Most common misuse.** *Confirmation sweep* — evidence gathered only in support of the frame, with no recorded attempt to find what would break it.

**When to adapt or skip** (SHOULD NOT skip; MAY adapt). D2 **SHOULD NOT** be skipped: a statement composed without a sourced evidence set has nothing for D4's challengers to trace. Quantitative and qualitative evidence **SHOULD** be paired where a clause asserts anything about behaviour, because a number with no account of why it moved, or an account with no number, each fail a different challenge later.

**Illustration (informative).** A subscription operator classifies its churn rate as proprietor-supplied, sourced to its billing records, and a competitor's published pricing as externally observable. It searches, and records, whether a recent price change rather than the suspected cause better explains the movement — the search that would disconfirm the frame, run and written down whatever it finds.

## 5.3 D3 — Compose



*Attribution.* This method's own transform, not adapted from a single named framework. The packaging convention it **MAY** hand downstream — a situation–complication–question form — is Minto's (1996); the composing rule itself is not.

**Purpose and position.** D3 turns a sourced evidence set and its frame into one candidate sentence: what is happening, to whom, since when and at what magnitude, with every clause resolved to the evidence that earned it and success stated as a baseline, a target and a date. It is the transform that produces the statement — the only artefact this method hands to Chapter 2.




**Figure 5.3 — The clause-to-evidence map.** A clause with no edge into the evidence column has not been earned. This method's format.

### Correct use.


- Every clause **MUST** resolve to at least one evidence item in the map, and **MUST NOT** assert more than its items support.
- Success criteria **MUST** be numeric and **MUST** state a baseline, a target and a date; an undated target cannot later be shown false. Where no one will propose a target at all, the verdict is **DEFICIENT** against the target field. A delegate's proposal is not an exception that lets a statement through with no commitment behind it: it is a commitment of a lesser kind, and it caps what D4 can later find — the verdict cannot rise above **PROVISIONAL** until the named decision-maker commits in their own right.
- The statement **MUST** restate its scope, so it can be used downstream without any further artefact.
- The statement **SHOULD** meet the six quality criteria Chapter 2's entry contract sets and **SHOULD** be a single sentence; D3 **MUST** record, for each, a pass or the reason for departure.
- The clause-to-evidence map **MUST** be retained, and **MUST NOT** be presented as part of the chain's derivation graph — it is Chapter 1's own artefact.

**Most common misuse.** *Target from nowhere* — a numeric target stated with no named person, delegate or decision-maker, standing behind it.

**When to adapt or skip** (SHOULD NOT skip; MAY adapt). D3 SHOULD NOT be skipped: without it there is no single sentence for D4 to test. The statement MAY be accompanied by a packaging form such as a situation–complication–question framing; the packaging carries no requirement of its own.

**Illustration (informative).** An operations director cannot yet commit to a target while a budget round is open; a named finance delegate proposes one, marked as proposed. D3 admits the statement on that footing, and the verdict that follows can go no further than PROVISIONAL until the director commits in person.

## 5.4 D4 — Test



**Attribution.** This method's own transform, not adapted from a single named framework. The independence it demands of a challenger is Nickerson's (1998) reason, applied here as the argument for a separate tester rather than a better author.

**Purpose and position.** D3 composes one candidate; D4 decides whether it survives being attacked. It assembles challengers with recorded standing, runs all nine challenges against the statement, records each outcome and issues exactly one verdict from the fixed vocabulary. It is the transform that enforces the falsifiability this chapter's title names, and the only transform whose entire output is a decision rather than a fact.

| Challenge                 | Challenger and standing          | Question        | Evidence cited | Outcome                                |
|---------------------------|----------------------------------|-----------------|----------------|----------------------------------------|
| 1 Falsification           | named challenger standing stated | question as put | evidence cited | withstood / forced a change unanswered |
| 2 Provenance              | named challenger standing stated | question as put | evidence cited | withstood / forced a change unanswered |
| 3 Baseline and commitment | named challenger standing stated | question as put | evidence cited | withstood / forced a change unanswered |
| 4 Symptom                 | named challenger standing stated | question as put | evidence cited | withstood / forced a change unanswered |
| 5 Ownership               | named challenger standing stated | question as put | evidence cited | withstood / forced a change unanswered |
| 6 Boundary                | named challenger standing stated | question as put | evidence cited | withstood / forced a change unanswered |
| 7 Prospective failure     | named challenger standing stated | question as put | evidence cited | withstood / forced a change unanswered |
| 8 Metric fitness          | named challenger standing stated | question as put | evidence cited | withstood / forced a change unanswered |
| 9 Confound                | named challenger standing stated | question as put | evidence cited | withstood / forced a change unanswered |

**Figure 5.4 — The challenge record.** Every row names who asked, on what standing, and how it was settled. This method's format.

### Correct use.

- All nine challenges MUST be run and recorded, each with its outcome and the challenger who put it; an unanswered challenge MUST NOT be recorded as withstood.
- Each challenger's standing MUST be recorded: what makes them competent to attack this specific claim, never seniority alone. Standing requires at least one challenger who is neither the composer nor the named decision-maker and who is accountable for the baseline metric, plus — where the statement asserts a monetary

consequence — one accountable for that consequence. Where no challenger with standing can be found, the verdict **MUST** be **DEFICIENT**, reason "no challenger with standing".

- Challenge and ratification are distinct acts with opposite incentives and **MUST** be performed by different people; the set **MUST NOT** be run as an approval body.
- **ADMITTED** **MUST** require a human challenger of record, however the challenge set was prepared or run.
- A **REFRAME** **MUST** name the evidence item that forced it; the superseded framing **MUST** be recorded and **MUST NOT** be re-adopted, and the reframe history **MUST** be acyclic.

**Most common misuse.** *Approval round* — the set run as a sign-off, with the challenger also acting as ratifier of the same statement.

**When to adapt or skip** (**MUST NOT** skip; challenges **MAY** be added, never dropped). D4 **MUST NOT** be skipped: the verdict is its only output, and no diagnosis reaches Chapter 2 without one. Engagement-specific challenges **MAY** be added to a diagnosis's own record; none of the nine published here **MAY** be dropped from it.

**Illustration (informative).** A statement claiming a retention gap is put to a challenger accountable for the retention measure, who asks what would show the claim false and finds the search was never run; the verdict names the missing search and holds at **DEFICIENT** until it is.

## 5.5 The nine challenges

The nine challenges attack a composed statement from four directions, over a spine of five checks that must simply hold. The spine tests falsifiability, provenance, the baseline and its commitment, ownership and prospective failure; the four attacks test the statement's level, its boundary, its measure and its explanation. Numbering follows the order the challenges entered the set, not the order a challenger would ask them, and does not change, because the challenge record cites by number.

**1. Falsification.** "What could I go and observe next week that would show this is false?" Pass condition, **MUST**: a named, obtainable observation that would show the statement false exists, and has been sought. A statement immune to any future observation has not been diagnosed; it has been asserted. Failing this challenge forces **DEFICIENT** — there is nothing yet for the chain to test against reality.

**2. Provenance.** "Take this clause. Which evidence item earned it, and from what system?" The question asks where a number came from, never what produced it. Pass condition, **MUST**: any clause the challenger picks resolves to a named source. A clause that traces to memory rather than a cited item forces **DEFICIENT**.

**3. Baseline and commitment.** "Where does the before number come from, and who signed up to the after number, by when?" Pass condition, **MUST**: the before number **MUST** name its system; the after number **MUST** name who committed to it and by when. Where the decision-maker cannot yet commit, a named delegate's proposal, marked as such, survives at **PROVISIONAL**; where no target exists at all, the verdict is **DEFICIENT**. This is where the fact of ratification lives — not as a separate record, but as this challenge's own pass condition.

**4. Symptom.** "What would have to be true for this to be the shadow of a different problem — does your evidence rule that out?" Pass condition, **SHOULD**: the diagnosis shows the evidence does not better support a superordinate problem. An unanswered or failed instance still forces a verdict: **REFRAME**, naming the evidence item responsible and the framing it supersedes.

**5. Ownership.** "Who acts on this, and can they act inside the scope you have drawn?" Pass condition, **MUST**: the named decision-maker holds authority within the stated scope. A statement whose named actor cannot in fact move within its own boundary forces **DEFICIENT** — a role has been named with no power behind it.

**6. Boundary.** "Name one thing you put out of scope that a reasonable reader would put in. Why is it out?" Pass condition, **SHOULD**: the exclusion is named and defended. Failing it forces **DEFICIENT** or, where the excluded item turns out to carry the real problem, **REFRAME**.

Challenge 4 tests the vertical axis: whether this is the right level of problem. Challenge 6 tests the horizontal axis: which adjacent matter was excluded, and on what reason. The two fail differently. Challenge 4 fails when a superordinate problem better explains the same evidence, and its remedy is a reframing. Challenge 6 fails when no defensible boundary can be given for an adjacent exclusion, and its remedy is usually a widened scope or a recorded reason, not a new framing. Merging them would let a team answer on one axis while failing the other unnoticed.

**7. Prospective failure.** "It is months on. The work that followed this statement solved the wrong problem. What was the statement wrong about?" Pass condition, **SHOULD**: every answer is either evidenced or recorded as a named gap. This challenge trades the tense of the others for a future one, forcing the composer to answer from inside an imagined failure rather than in the statement's own defence. Klein's (2007) project premortem asks a team to imagine the failure before it happens. Mitchell, Russo and Pennington (1989) found, in laboratory tasks, that placing an outcome in an already-

happened past rather than an open future changed the content of the explanations people gave for it — a directional laboratory finding about explanation, not evidence that the technique improves what a diagnosis produces. Failing this challenge forces REFRAME or DEFICIENT.

**8. Metric fitness.** "Is that the right measure of the thing you say is wrong? I would have looked at another — why is yours the one that settles this?" Pass condition, MUST: the headline metric is defended against at least one named alternative, or changed. A challenger already required to be accountable for the baseline metric cannot be handed an optional question about it; failing this challenge forces REFRAME or DEFICIENT.

**9. Confound.** "What else could have produced this movement — a season, a price change, a change in how you measure it, a reorganisation? What rules those out?" Pass condition, SHOULD: each named alternative is ruled out by a cited evidence item or recorded as a named gap. Failing it forces REFRAME or DEFICIENT.

Challenge 9 is not challenge 1 asked twice. Challenge 1 asks for a future observation not yet obtained; challenge 9 asks the composer to rule out an explanation the evidence already contains. Challenge 1's failure is an unfalsifiable statement; challenge 9's is a spurious one — real numbers attached to the wrong story. Folding the two together would let a statement pass on one clean falsifier while a confound sat unexamined underneath it.

The set is fixed at nine. A diagnosis MAY record engagement-specific challenges of its own alongside these, but the published set MUST NOT grow beyond nine without a new changelog entry recording the change: an adversarial set that keeps growing becomes the checklist this stage exists to replace.

| B | CHALLENGE                 | ADMITTED | PROVISIONAL | DEFICIENT | REFRAME |
|---|---------------------------|----------|-------------|-----------|---------|
|   | 1 Falsification           |          |             | ✓         |         |
|   | 2 Provenance              |          |             | ✓         |         |
| ✓ | 3 Baseline and commitment |          | ✓           | ✓         |         |
|   | 4 Symptom                 |          |             |           | ✓       |
|   | 5 Ownership               |          |             | ✓         |         |
|   | 6 Boundary                |          |             | ✓         | ✓       |
|   | 7 Prospective failure     |          |             | ✓         | ✓       |
| ✓ | 8 Metric fitness          |          |             | ✓         | ✓       |
|   | 9 Confound                |          |             | ✓         | ✓       |

A tick in the B column marks a challenge whose pass condition names the baseline metric.

**Figure 5.5 — The nine challenges and the verdicts they can force.** Some can only be put by a challenger accountable for what the statement measures or what it costs. This method's format.

## 6. Evidence, provenance and the handoff

### 6.1 Who can supply a fact

A fact that will earn a clause of the statement comes from one of two places, and D1's evidence-demand list MUST classify every item accordingly: proprietor-supplied, meaning only the organisation itself holds it — an internal record, a report nobody outside sees, a decision already taken — or externally observable, meaning any competent party could in principle obtain it — a published index, a regulator's filing, a competitor's own disclosure.

The classification is not bookkeeping; it decides what a missing item means. A missing externally observable fact is recorded as an open research item and does not halt D2 — the fact goes on existing whether or not anyone has fetched it yet. A missing proprietor-supplied fact is different: if the organisation itself cannot produce it, no further search will, and the gap MUST name which §3.2 field it leaves unevidenced. D3's admission check then refuses a statement composed over that gap.

This is also what makes a statement valid for *this* organisation rather than a generic one. An externally observable fact could belong to any comparable organisation in the same position; a proprietor-supplied fact — the baseline metric, the decision-maker's own account of scope, the record a number came from — is what anchors the statement to the

situation actually in front of the named decision-maker, not to an industry average standing in for it. A frame built entirely on externally observable facts has described a market; it has not yet described a problem this organisation owns.

**Illustration (informative).** A logistics operator's fuel-cost trend is externally observable: several data providers track it. Its own on-time delivery rate against a contracted service level is proprietor-supplied — no outside source holds it, and its absence is not a research gap but a missing frame.

## 6.2 Sourcing and provenance

Chapter 2 cites a finding by its code, and a code's derivation chip shows where that finding came from. The diagnosis stage **MUST NOT** mint codes, so nothing produced here is cited that way. Every evidence item **MUST** instead name its source directly — the record, the document, the person accountable for the number — and **SHOULD** be dated; that named source, not a code, is what a later reader or a challenger follows back.

"Cited by provenance" means exactly this: a clause traces to a source stated plainly in the evidence set, not to an identifier standing in for one. A baseline metric **MUST** exist with a named source, or the gap is recorded and instrumentation is named as the first work still to do. Context sufficient for the frame's clauses **MUST** be gathered, and **SHOULD** go no further — the evidence-demand list, not a general appetite for research, decides when D2 is finished. D2 **SHOULD** also record a deliberate search for evidence that would disconfirm the frame, whatever it finds; an evidence set that only confirms is a sweep, not a search.

## 6.3 The internal-fact hook

One kind of proprietor-supplied fact has a second option beyond a named source. Internal facts gathered at D2 **MAY** be registered as Chapter 2's own internal-fact codes,  $I \cdot 1 \dots I \cdot n$  (Chapter 2, §6.1 and Appendix A); where they are not, they are cited by provenance as §6.2 describes, and the source **MUST** still be named. The level stays **MAY**, the family stays Chapter 2's, and Chapter 2's own minting window — before grading, at stop 4 — already reaches back far enough to cover a fact gathered during diagnosis, so nothing about this chapter's timing requires any change on Chapter 2's side.

It would be easy to mistake this for a second interface between the chapters, alongside the statement. It is not: **the problem statement is the only connection between the chapters**. Registering a fact as an I-code is an act Chapter 2 performs on Chapter 2's own registry, under Chapter 2's own rule, on a fact the organisation already held before D1 ever opened a frame around it. Nothing this chapter composes, tests or reaches a verdict on travels along that registration — no clause, no challenge outcome, no verdict. A diagnosis that reaches **ADMITTED** without a single fact ever registered as an I-code has lost nothing; the hook is a convenience for a fact that will be reused later, not a requirement this chapter imposes or a channel it opens.

## 6.4 Where the diagnosis trail ends and the chain's graph begins

Chapter 2 already states the rule that governs this boundary: every citation path **MUST** end at a root backed by a named external source, a registered internal fact or the entry problem statement (Chapter 2, §6.4). The entry problem statement is one of those roots already, and nothing here adds to that list or subtracts from it. Because no transform mints a code (§6.2), the diagnosis trail — the evidence set, the clause-to-evidence map, the challenge record — has nowhere to travel into that graph, and does not try to. It terminates at the statement and **MUST** be retained alongside it; the statement is what stop 1 receives.

As §6.3 states, there is no side channel here and no second interface: the statement is the only connection this method defines. The one exception a careless reading — or a careless figure — could mistake for a second connection is the internal-fact hook. An organisation's own fact **MAY** be registered under Chapter 2's own  $I \cdot 1 \dots I \cdot n$  family, at Chapter 2's own **MAY** level. That registration is not something this chapter does or produces; it is Chapter 2's own act, on Chapter 2's own artefact, performed on a fact the organisation held before the diagnosis ever began. Nothing this chapter's evidence set, clause-to-evidence map or challenge record adds to that fact travels along the  $I \cdot i$  edge. The

edge rises from the organisation's own internal facts into the chain's root layer; it never leaves this chapter's evidence set. Reading it the other way round — as an edge departing from what D1 through D4 compose — would put a Chapter 1 artefact inside Chapter 2's graph, which is exactly what the rule stated above forbids.




**Figure 6.1 — Where the diagnosis trail ends and the chain's graph begins.** The statement is the only connection: the chain cites it as a root and never reaches past it. This method's format.

## 7. The challenge set, failure modes, and what this paper does not claim

### 7.1 The bar

A statement can satisfy every criterion the entry contract publishes and still be a well-argued case for the wrong problem. Specific, quantified, evidence-based, solution-free, scoped, anchored to a stakeholder, about what is happening rather than why — each of the six criteria Chapter 2's entry contract lists is satisfiable by a statement built carefully around a symptom. A checklist asks a statement about its own shape; it has no way to ask whether that shape sits over the right problem.

The bar this stage sets is therefore adversarial: a statement is admitted only once it has been put in front of named people whose task is to find what would break it, and has survived. The person who composed the frame is the wrong person to run that search. Nickerson's (1998) review of confirmation bias found reasoners repeatedly better at defending a belief than at attacking one — a finding about reasoning in general, extended here, as this chapter's own analogical step, to the person who wrote the statement being tested.

What follows specifies who may run the attack and what happens when a statement does not survive one; the nine forms it takes are set out in §5.5. A statement that has withstood every challenge is challengeable, not thereby correct: it has been put in front of people equipped to break it, and did not break.

### 7.2 Who may challenge

Competence to challenge is fixed by what the statement claims, not by who holds the more senior chair. A statement asserting a fall in retention is challenged by whoever answers for the retention measure; a statement asserting a cost is challenged, in addition, by whoever answers for that cost. A senior sign-off offered without either accountability tests nothing the statement actually says.

At least one challenger **MUST** be neither the statement's composer nor its named decision-maker, and **MUST** be accountable for the baseline metric the statement moves. The decision-maker is excluded because they commissioned the framing and stand to gain or lose by its verdict, so they cannot also be the test of it. Composer and decision-maker may answer questions and supply evidence; neither may be the one putting them.

Where the statement asserts a monetary consequence, a challenger accountable for that consequence — in practice, someone in finance — **MUST** also take part. The requirement is conditional on the claim, never a standing seat: a statement asserting no cost gives a finance challenger nothing to falsify, and seating one anyway would be exactly the ritual this design exists to prevent.

Each challenger's standing **MUST** be recorded — what makes this person, and not another, competent to attack this specific claim, never seniority alone. A list of names with nothing said about why any of them could break the statement is a signature round wearing the challenge record's shape. Where no challenger with standing can be found, the verdict **MUST** be **DEFICIENT**: the accurate result is a diagnosis the stage refuses, not a challenge run by whoever happened to be free.

Challenge and ratification are not the same act performed twice, however similar the room looks. A challenger succeeds by finding a flaw in the statement; ratification succeeds by the statement standing — opposite incentives, and one person cannot perform both without one swallowing the other. The nine challenges **MUST NOT** be run as a sign-off round in which everyone present already agrees; a set run that way answers no question it was built to ask.

**ADMITTED** **MUST** require a human challenger of record, however the challenge set was prepared or run, and a diagnosis **MUST NOT** reach it without one. This is a property of the verdict, not an aside about how a challenge set might come to exist.

### 7.3 The nine challenges

The nine challenges now sit at §5.5, immediately after D4, the transform they test.

### 7.4 The challenge record

The challenge record is Chapter 1's audit artefact — the document a reviewer reads first, and the analogue of Chapter 2's code registry for a stage that mints no codes of its own. It carries one row per challenge, nine rows always, each naming the challenger, the question put, the answer, the evidence cited and the outcome; an unanswered challenge is never recorded as withstood.

Retention is the record's load-bearing property, not a filing habit. Where a challenge forces **REFRAME**, the superseded framing and the record of the challenge that forced it out are both kept, never discarded once the new framing is composed. A reframe history that could be edited after the fact would let anyone claim it was evidence-forced and acyclic without a later reader having any way to check either claim. Keeping the losing framing is what makes a **REFRAME** a finding rather than a second draft.

The challenge record **MUST** be retained with the statement, including the records of any superseded framings: a statement kept apart from its record is, for the purposes of this stage, unaudited, whatever verdict it once carried.

### 7.5 Failure modes

Each transform's most common misuse is a failure mode its own checks exist to catch. One per transform:

1. **D1 — Solution in disguise.** A frame that already carries its own remedy inside the wording of the situation, so that describing the problem and choosing the fix become the same sentence, and no evidence can any longer disconfirm either one.
2. **D2 — Confirmation sweep.** A search that only ever finds what the frame already expected, because no one went looking for the fact that would have broken it.
3. **D3 — Target from nowhere.** A numeric target with a baseline and a date but no name behind it: nobody committed to it, so nothing in the statement says whose responsibility it is when the date arrives.
4. **D4 — Approval round.** The nine challenges run as a meeting where everyone present already agrees: the people who framed and composed the statement nod it through rather than attacking it, and whoever else is in the room follows their lead. A challenger who wants the statement to survive is not challenging it. A diagnosis run this way produces a record with nine filled rows and no test behind any of them. The remedy is not a better-run meeting, but putting the statement in front of someone with standing to attack the specific thing it claims, and letting them try.

### 7.6 Audit procedure

An audit of a diagnosis reads backwards from its verdict, the same direction Chapter 2's own audit reads backwards from a pillar.

Start with the verdict and what it names. For **ADMITTED**, that is the challengers and their standing, and who committed to the target and by when. For **PROVISIONAL**, it is the missing act. For **DEFICIENT**, it is every field left unevidenced, or the absence of a challenger with standing. For **REFRAME**, it is the evidence item that forced it and the framing it replaced.

Move to the challenge record next. Confirm all nine rows are present, each naming a challenger and an outcome, and that no unanswered challenge is recorded as withstood. Where a challenge names a gap rather than a source, confirm the gap is named against a specific field, not left general.

From there to the clause-to-evidence map: walk every clause of the statement to the item or items that earn it, and confirm none is asked to carry more than its evidence supports.

Finally to the sources themselves: each named, each dated where the record claims a date, each still saying what the map claims it says.

A diagnosis that survives this reading has been tested against attack; it has not been declared true.

## 7.7 What this paper does not claim

**No ownership of the named techniques.** The falsification question and the pre-mortem framing (Klein, 2007) are general practice, not this chapter's invention; goal-setting language belongs to Locke and Latham. This chapter claims the two-sided contracts, the challenger model, the nine challenges, the verdict taxonomy and the reframe rule, and nothing beyond them.

**No causal link from diagnosis to outcome.** No study isolates diagnosis quality as an independent variable against a business outcome, and this chapter does not claim to have found one. Wedell-Wedellsborg (2017) measures executives' perception of a costly diagnosis, not the cost itself. Nutt (1999) links decision tactics to whether decisions were carried through successfully, correlationally, across 358 decisions, not diagnosis quality specifically and not by experiment. Lyles and Mitroff (1980) study the formulation process itself, not what it produces downstream. Locke and Latham (2002) and Nickerson (1998) anchor two ingredients of the method in domains that are not strategic diagnosis. None of this shows that a diagnosis produced this way leads anywhere better than one produced another way.

**No benchmark against another diagnostic method.** Nothing here compares this decomposition with a linear root-cause chain, a design-thinking framing stage or an unstructured problem conversation on outcome quality.

**No claim that a competent challenger is a correct one.** The challenger model bets that accountability for a measure produces sharper attacks on claims about that measure. It is a reasoned bet, not a measured one, and no source in this chapter tests it.

**Analogical extensions this chapter owns, not its sources.** Card (2017) is a peer-reviewed critique of root-cause analysis in healthcare; its application against strategic diagnosis is this chapter's own extension, not a finding of the source. Nickerson (1998) studies reasoning in general, not organisational challenge. Ordóñez and colleagues (2009) argue that goal setting carries systematic side effects; Locke and Latham (2009) dispute that argument in the same issue, and this chapter cites the dispute as live, never Ordóñez alone as settled.

**Withstanding challenge is not truth.** A statement that survives the nine challenges has been put under attack and answered with evidence. It has not thereby been shown correct.

**Not advice, and free to use.** This chapter describes a method; it is not legal, financial or investment advice, and may be used freely by anyone.

## 8. Applicability

Four transforms and nine challenges cost more than a paragraph that states a problem and calls it done; that cost buys something specific, and it is not owed to every concern that reaches a decision-maker's desk.

### Use it when

- A statement will be tested by people who did not write it, and the record of that testing has to survive a later question about why the frame was drawn where it was.
- A diagnosis of the wrong problem, however cleanly solved afterwards, is expensive twice over (Rumelt, 2011a).
- More than one function has a stake in the answer — the metric owner, whoever answers for a monetary consequence, the person who will act on the finding — and each would supply a different fact if asked alone.
- There is time to search deliberately for what would disconfirm the frame and to convene challengers with standing before a target is committed.

### Lighten it or choose another approach when

- No one who is neither the composer nor the named decision-maker can be found or freed to sit as a challenger. Running the set thin, with names but no standing, produces a signature round, not a diagnosis, and is worse than naming the gap and stopping.
- A disconfirming search and independently sourced context cannot be assembled without gaps; a partial pass, with those gaps named, is more defensible than a full diagnosis run on thin evidence.

## Preconditions

Before starting the first transform, confirm:

1. a decision-maker who can be named by role, with a stated concern and something that has changed;
2. access to evidence from more than one source;
3. at least one prospective challenger who is neither the composer nor the named decision-maker, and who is accountable for a metric the concern touches;
4. time for four transforms and nine challenges before any target is committed, which a crisis timeline does not allow.

A decision-maker who wants an answer by Friday and a diagnosis that clears every precondition are not always compatible. The transparent response to that tension is to say so before starting, and to choose the lighter path deliberately, not to clear a precondition quietly and call the record complete.

## 9. Why this decomposition and not the alternatives

Four transforms are not the only way to cut diagnosis into pieces. Four candidates were considered instead: a single flat contract, a longer sequence of smaller steps, folding composition into testing, and folding framing into evidence-gathering; each is rejected for a stated reason, not on preference.

**Against a flat contract.** The analysis chain's own entry contract already states what a statement must contain: five prerequisites, sourced evidence for every clause, six quality criteria, a single-sentence form (Chapter 2, §3.2). Restating that list is the one thing this chapter must not do. A flat contract could not do what four transforms add, in any case. It cannot express an admission check that refuses an unfit input rather than transforming it regardless; it cannot locate a deficiency at a specific field rather than announcing it against the whole. It cannot distinguish "the evidence is missing" from "the statement is wrong" — the two call for different remedies. Nor can it attach a requirement, a human challenger of record, to one transform's admission check and to one verdict, rather than to a document considered as a whole.

**Against a longer chain.** A longer chain is not merely imaginable. An earlier decomposition of the same problem, predating this method, ran to eight stages instead of four. A structural comparison checked each stage against a fixed input, a fixed output artefact and a scored contract check. Six of the eight carried something resembling a check, but softer than the standard this chapter sets: a self-reported sufficiency review, a required diagram, an exact-vocabulary verdict. Two of the eight were not transforms at all. One selected among a fixed set of lenses without producing a decision-relevant artefact of its own. The other fanned a single input out to several parallel gatherers who reported back to the same place, adding a box without adding a check. None of the eight cleared a fixed input, a fixed output and a scored check together. Two further candidate steps sit outside the exit contract entirely. One is an account of why the situation arose, which the entry contract already excludes from the statement. The other is a rehearsal of how the diagnosis might later prove wrong, whose substance this chapter keeps as one of the nine challenges, without giving it a transform of its own. Remove the steps that only route or fan out, and tighten every remaining check to the two-sided standard this chapter sets. Four remain: frame, evidence, compose, test.

**Against merging composition into testing.** Composing a statement and testing one need different things admitted. Composition needs a closed evidence set with no gap against a required field; testing needs challengers with recorded standing, and specifically excludes the person who composed the statement. Merging the two would seat the composer in the challenger's chair — a person least placed to look for what would break their own frame (Nickerson, 1998). The rule against it is not a suggestion.

**Against merging framing into evidence-gathering.** The frame's evidence-demand list is what makes evidence-gathering a finite task rather than an open search; it is the stopping rule, and the thing the proprietor-supplied and externally-observable classes classify. Collapse the two transforms, and "minimal market and customer context" no longer has anything that defines minimal.

This decomposition is reasoned, not empirically benchmarked. The earlier decomposition cited above is this method's own prior design work, not an independent study. No source compares this ordering, or any ordering, against another on outcomes. The case for four transforms rests on what each admission check and each contract check is shown above to add, not on a measured result.

## Extended rationale: foundations and contribution

Lyles and Mitroff (1980) describe the formulation process itself, not evidence that careful formulation produces a better result. Rumelt's 2011 *McKinsey Quarterly* article opens the kernel of good strategy with diagnosis, naming the challenge before any guiding policy; Wedell-Wedellsborg (2017), surveying 106 executives, found most moved into solution mode before checking their own understanding — a perception survey, not a measured result. Nickerson's (1998) review of confirmation bias found reasoners seek evidence fitting an existing belief; applied here, whoever names a frame is

poorly placed to find what would break it — an extension this chapter makes, not a strategy-specific finding. Locke and Latham (2002) found specific, difficult goals raise task performance, not diagnosis quality; Ordóñez et al. (2009) argued prescribing goals too freely carries side effects, which Locke and Latham (2009) disputed directly — this chapter states the dispute as live.

None of this chapter's material is new: a five-field intake, a named-owner check and a rule against smuggling a cause into a symptom's description already circulate as practical craft. Its contribution is the two-sided contract at each transform, a challenger model excluding the statement's own composer, a fixed nine-challenge set, a closed-vocabulary verdict, and an acyclic reframe rule — kept once, not claimed as invention.

## 10. Scope, roadmap and references

### 10.1 Scope of Chapter 1

Chapter 1 specifies the diagnosis stage: the two-sided contract of each transform — Frame, Evidence, Compose and Test — the challenger model, the nine challenges and the four verdicts, every rule marked by level. It is written as a reference against which a diagnosis can be run and then audited. It states, transform by transform, what each admits, what it refuses, what it produces, and the condition its output must meet before the next transform, or the analysis chain, may rely on it.

This methodology sits here; the worksheets, question scripts, scoring and a full worked case that put it into practice sit in the practitioner guide.

It does not cover:

- the analysis chain that follows an admitted statement, whose own entry contract this chapter's statement is built to satisfy (Chapter 2, §3.2);
- what happens once a target is committed — mobilisation, execution and benefits realisation;
- evidence that four transforms outperform another decomposition: the sequence is reasoned, not empirically benchmarked (§9).

The four verdicts, the challenger model and the reframe rule are stated here in full; none of them is redefined, weakened or repeated anywhere else in the series.

Illustrations throughout are generic and informative: they set no requirement, name no company and fix no count, except the method's own fixed sets — four transforms, four verdicts, nine challenges.

### 10.2 Roadmap

The following are planned, with timing set by the author's prioritisation, each keeping the rule "No output without a traceable input".

- **Practitioner guide** — worksheets, question scripts, a scoring rubric and a full worked case for the four transforms and the nine challenges, in preparation.
- **A later revision of this chapter** — the nine challenges are the ceiling as this chapter stands; a tenth would need its own evidence, not a preference, and would be recorded as a new changelog entry, not folded into this one.

This fills the slot Chapter 2's own roadmap already reserves for the diagnosis stage. The chapters that hand off from the chain's pillars remain Chapter 2's to schedule.

### 10.3 Changelog

| Date              | Change                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 22 September 2026 | <b>First published.</b> Four transforms — Frame, Evidence, Compose, Test — each with its own admission check and contract check, the challenger model, nine challenges, the four-verdict taxonomy, and the reframe rule.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| 28 September 2026 | <b>Restructured for reading order, no rule changed.</b> The nine challenges (§7.3) moved to §5.5, immediately after D4; §4.3 gained a "Key rules" column absorbing Appendix A's content, and Appendix A was removed; §2 Foundations and §3.4 were condensed to a front residual, with the fuller argument moved into §9's new "Extended rationale"; §1.1, §1.2 (explanatory prose only, not the RFC boilerplate or any term definition), §1.3, §3.1, §3.2, §3.3, §4.1, §4.2, §4.4, §7.4, §8, the Abstract and §4.3's internal-fact-hook mention were condensed or tightened. No transform, admission check, contract check, challenger rule, challenge or verdict was added, removed or reworded, and every rule sentence keeps its original wording. |

| Date              | Change                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 28 September 2026 | <b>§1.2's RFC 2119/8174 boilerplate replaced with a pointer to the shared conventions page (<a href="http://whitepaper.strategyascode.org/conventions/">whitepaper.strategyascode.org/conventions/</a>); no rule reworded, no MUST/SHOULD/MAY meaning changed, only relocated.</b> §1.2's own chapter-specific vocabulary (the gate/stop note, the transform notation, the term glossary) is unchanged. |
| 28 September 2026 | <b>§7.1 wording corrected to point to the challenge set's new location (§5.5); no rule changed.</b>                                                                                                                                                                                                                                                                                                     |
| 28 September 2026 | <b>Version numbers removed; changes are now recorded by date; no rule changed.</b>                                                                                                                                                                                                                                                                                                                      |

This changelog belongs to Chapter 1 alone; it leaves Chapter 2's own changelog untouched. A future change to this chapter is recorded as a new entry here, never as a silent edit to this one.

## 10.4 References

- 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>
- A. J. Card, "The Problem with '5 Whys'" (*BMJ Quality & Safety* 26(8), 2017) — <https://doi.org/10.1136/bmjqs-2016-005849>
- 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>
- E. A. Locke and G. P. Latham, "Building a Practically Useful Theory of Goal Setting and Task Motivation: A 35-Year Odyssey" (*American Psychologist* 57(9), 2002) — <https://doi.org/10.1037/0003-066X.57.9.705>
- E. A. Locke and G. P. Latham, "Has Goal Setting Gone Wild, or Have Its Attackers Abandoned Good Scholarship?" (*Academy of Management Perspectives* 23(1), 2009) — <https://doi.org/10.5465/AMP.2009.37008000>
- M. A. Lyles and I. I. Mitroff, "Organizational Problem Formulation: An Empirical Study" (*Administrative Science Quarterly* 25(1), 1980) — <https://www.jstor.org/stable/2392229>
- B. Minto, *The Minto Pyramid Principle: Logic in Writing, Thinking, and Problem Solving* (Minto Books International, 1996) — <https://www.barbaraminto.com/>
- 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>
- R. S. Nickerson, "Confirmation Bias: A Ubiquitous Phenomenon in Many Guises" (*Review of General Psychology* 2(2), 1998) — <https://doi.org/10.1037/1089-2680.2.2.175>
- P. C. Nutt, "Surprising but True: Half the Decisions in Organizations Fail" (*Academy of Management Perspectives* 13(4), 1999) — <https://doi.org/10.5465/AME.1999.2570556>
- L. Ordóñez, M. E. Schweitzer, A. D. Galinsky and M. H. Bazerman, "Goals Gone Wild: The Systematic Side Effects of Overprescribing Goal Setting" (*Academy of Management Perspectives* 23(1), 2009) — <https://doi.org/10.5465/AMP.2009.37007999>
- R. Rumelt, *Good Strategy/Bad Strategy* (Crown Business, 2011a) — <https://www.penguinrandomhouse.com/books/208668/good-strategy-bad-strategy-by-richard-rumelt/>
- R. Rumelt, "The Perils of Bad Strategy" (*McKinsey Quarterly*, 2011b) — <https://www.mckinsey.com/capabilities/strategy-and-corporate-finance/our-insights/the-perils-of-bad-strategy>
- T. Wedell-Wedellsborg, "Are You Solving the Right Problems?" (*Harvard Business Review*, January–February 2017) — <https://hbr.org/2017/01/are-you-solving-the-right-problems>

## 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](http://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: A Falsifiable Method for Strategic Diagnosis*, Chapter 1. <https://whitepaper.strategyascode.org/diagnosis/>. Chapters 2, 3 and 4 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.