

The AFS Casebook

Fitness Is Proven, Not Presumed

By Mario Aiello — Adaptive Ways

An illustrated field guide for recognising and testing presumed organisational fitness.

Contents

The argument is arranged by practical purpose rather than case number. Use the contents to follow the reasoning, the tools to run a live fitness check, and the Casebook Index for direct numerical lookup.

Opening	3
How to Use This Casebook	5
Part I — The Governing Law	6
Part II — The Four Questions for Fitness	10
Part III — From Questions to Design	15
Part IV — The Fitness Sequence	18
Part V — From Evidence to Action	23
The 20-Minute Fitness Check	27
Part VI — The FVI Lens	28
The Fitness Claim Record	34
Part VII — The Friction-Fitness Map	35
Lineage, Boundaries, and Debt	40
Closing — The AFS Paradox	41
Casebook Index	42

Opening

I still have the first slide deck I ever built for a client on agile transformation, and I keep it the way some people keep an old passport photo — mortifying, and useful precisely because it's mortifying. It promised a great deal. It used the word “framework” without irony, several times per page. I believed every word of it, which is the part that actually matters.

That deck was Simple Agility, though it didn't have the name yet — it was just what I did, before I'd noticed it was a thing that could have a name. Writing it down became The Agile Operating System, which was a more honest and considerably less comfortable exercise, the way most attempts to describe your own actual process tend to be. Beyond Agile came from watching organizations apply agile ceremonies to problems that had nothing to do with software delivery cadence and everything to do with structural misfit, and watching everyone conclude, reasonably, that agile “didn't work here.” It wasn't that agile didn't work. It was that agile was being asked a question it didn't have an answer to.

AFS is what happened when I stopped trying to fix agile's edges and asked what the underlying claim should have been all along: not “adopt these practices” but “prove, continuously, that whatever you've built still fits the conditions you're actually in.”

I want to be clear about what this book is and isn't. It isn't a sales document, and if that sentence itself reads like the opening of one, I understand the suspicion and I'd share it. It's an attempt to lay out an approach — never a framework, and that distinction matters more than it sounds like it should — as plainly and as sceptically as I can manage, including toward my own work.

Here's the other thing it is, and the reason it looks the way it does: somewhere around the fortieth client engagement, I noticed I kept meeting the same four people. Different companies, different industries, different countries, same four people, wearing different lanyards. There's always an Executive who wants a clear story and is prepared to believe it slightly before the evidence arrives. There's always a Guide, who has learned that the fastest way to move a room is to ask one short, quiet question rather than deliver a lecture. There's always a Team, doing the actual work, absorbing the gap between the plan and the Tuesday it lands on, usually without complaint and usually without being asked. And there's always Reality, who was never on the invite list, never RSVPs, and shows up anyway — in the support queue, in the exception log, in the customer who calls back twice.

This book is built around those four. Its purpose is narrower than explaining every part of AFS: it helps you recognise a claim of organisational fitness, test that claim against evidence, and decide what must be checked or changed next. The prose provides the reasoning. The forty-five illustrated scenes provide recognition. The practical pages turn recognition into a conversation you can run.

Read it in three ways. Read straight through for the argument. Enter through any cartoon when a scene feels familiar. Or use the tools to test a live claim in your own organisation. The cases are arranged by concept; the numerical index at the back gives you a second route through them.

Start with the law. Then make it earn its keep in a real room.

How to Use This Casebook

This is not an implementation manual and it is not a catalogue of answers. Use it as a disciplined interruption when a plan, design, process, or operating model is being treated as fit without anyone having stated what would prove otherwise.

Three useful entry points

In a meeting: choose the case that most resembles the current situation. Ask what claim the room is making, what evidence supports it, and what evidence would change the decision.

In a review: use the Four Questions to define the fitness claim, the FVI Lens to test its evidence, and the Friction-Fitness Map to locate pressure that the formal report may be hiding.

In practice: record one explicit claim, one expiry condition, and one next check. If the conversation produces only agreement, language, or a new diagram, it has not yet tested fitness.

A warning about the vocabulary

AFS is meant to interrogate whatever design is in use, not become another design to install. Its terms are useful only when they sharpen a question or change a decision. When the vocabulary becomes the performance, stop and return to the claim, the evidence, and Reality.

Part I — The Governing Law

There is a moment in almost every transformation program where someone stands up and says the word is now working. Not “working well.” Not “working for whom.” Just working — the way you'd say a car is running, before anyone asks whether it's running toward the cliff or away from it.

I have stood in that room. I have, on at least one occasion I'd rather not itemize, been the one who said it.

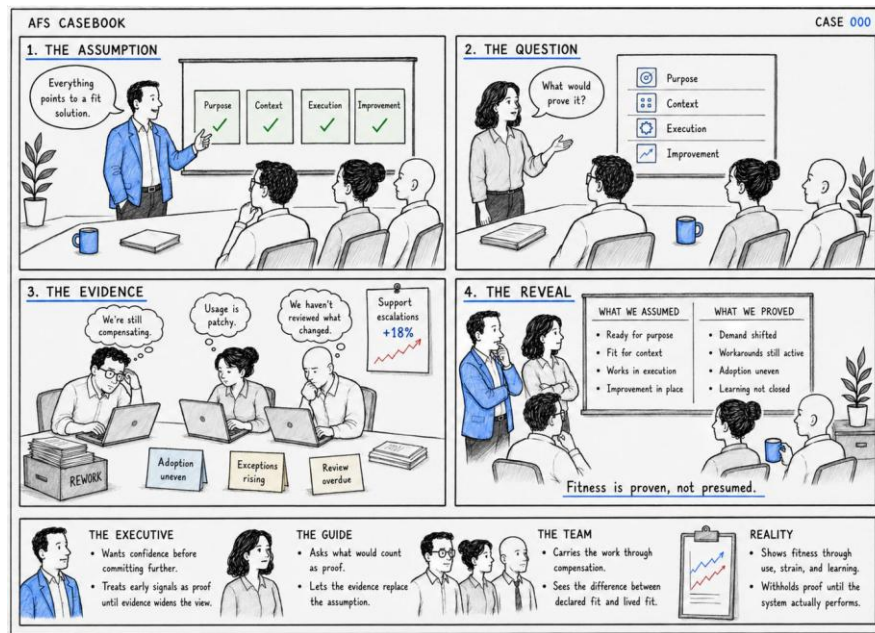
The trouble with “it's working” is that it's a claim about the present tense pretending to be a claim about the future. Nobody means to lie. The dashboards are green, the retro was upbeat, the sponsor is happy — for now. But “for now” is doing all the load-bearing work in that sentence, and nobody puts it in the deck.

This is the problem AFS exists to name, and naming it plainly is the whole first move: fitness is proven, not presumed.

Say it again slower, because it sounds like a slogan and it is not meant to be one. Fitness — whether an org design, a process, a strategy, a team structure fits the conditions it's actually operating in — is not something you get to assume because it was designed carefully, approved at the right level, or has worked before. It has to be demonstrated, continuously, against the conditions as they currently are. The moment you stop demonstrating it and start assuming it, you have left fitness behind and started managing a story about fitness. Organizations are extremely good at managing stories. It is, arguably, our core competency as a species that builds institutions.

This law comes from watching the same failure mode recur across three decades and several industries: a structure gets built to fit a specific context, it works, everyone concludes it works because it's well-designed, and then the context moves — the market shifts, the team turns over, the regulation changes — and the structure keeps being praised for the very features that no longer apply. The org chart survives the reason it existed. This happens in banks. It happens in insurance. I would wager it happens in your kitchen, in a system for loading the dishwasher engineered around a household that has since had a baby.

Here's the case that opens the whole casebook, because it's the cleanest statement of the law we ever drew.

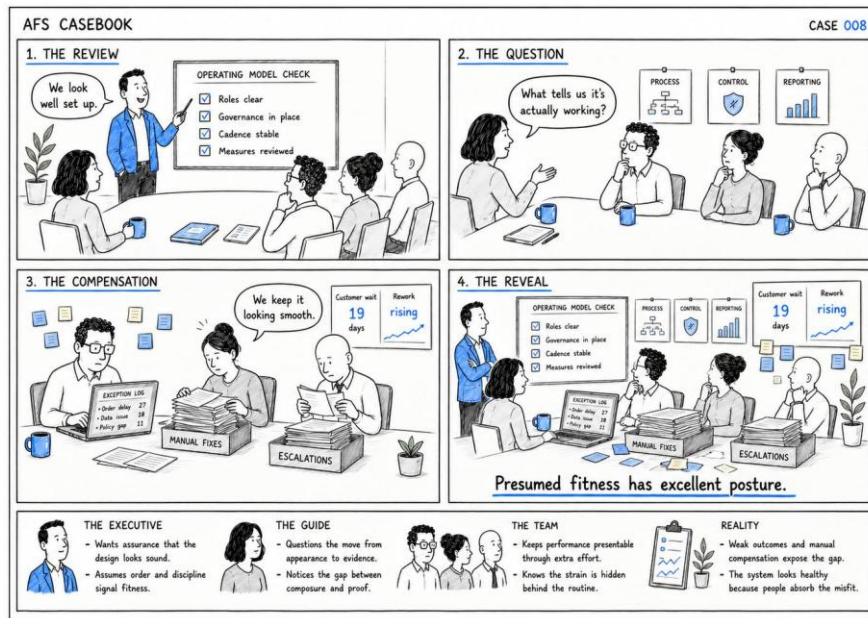


Case 000

FIELD NOTE — Claim: this design is fit. Test: fit for what conditions, and what would disprove it? Next move: name one observable failure condition before approving the claim.


Here is the self-criticism, up front, because AFS earns no exemption from its own law: a Governing Law that says “keep proving it” can itself become a thing people presume rather than practice. I have watched practitioners recite “fitness is proven, not presumed” in a tone identical to the one they used for the framework it replaced. The law does not protect itself. It only works if the organization actually does the proving.

Presumed fitness, meanwhile, has extremely good posture. It sits up straight in the steering committee. It has never once slouched.



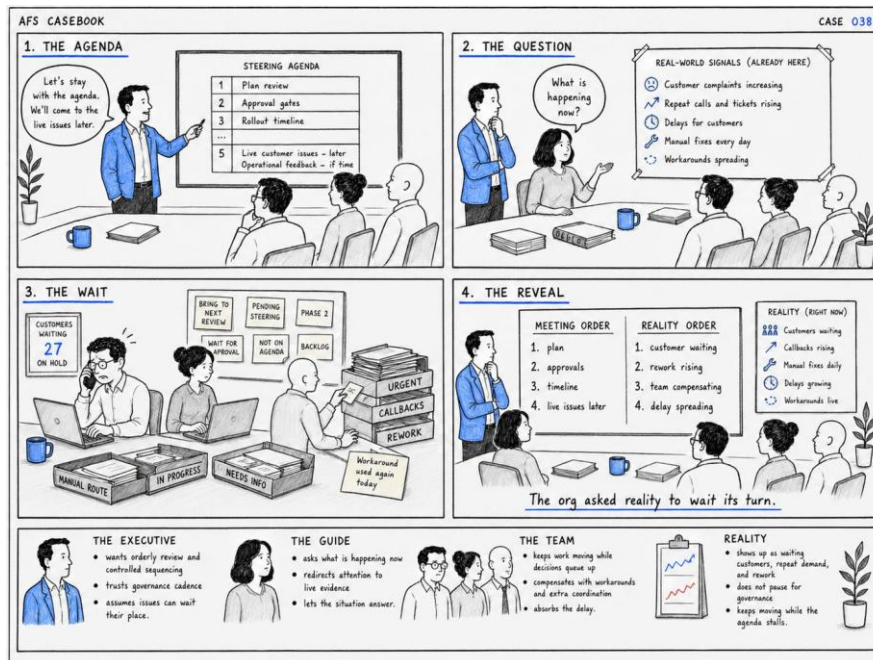
Case 008

A related habit: announcing a change and quietly treating the announcement itself as the change. It is a very efficient way to feel like something has moved. Reality, notably, does not attend announcement meetings. It was never on the invite list, and it has never once RSVP'd.




Case 001

What does the law demand in practice? One explicit fitness claim, one account of the conditions under which it is supposed to hold, and one statement of what evidence would make the organisation reconsider. Not paperwork. Not another audit function bolted onto the org chart. A live claim with a credible way to fail.



Case 038

There's also a quieter failure mode than presuming outright, and it's the one that shows up when everyone is, genuinely, trying. Effort goes up. Hours go up. Meetings held, tickets touched, tasks moved — all of it climbing on the chart. And the actual fitness of the system — the thing customers experience, the thing that was supposed to improve — does not move, or moves the wrong way, because effort and fitness are different axes and nobody drew the second one.



Case 032

The Governing Law is not clever. Its value is that it is inconvenient. It refuses to let “fit for purpose” become a certificate you earn once and frame on the wall. Before you can prove fitness, though, you need to know what you're actually asking. That's the Four Questions.

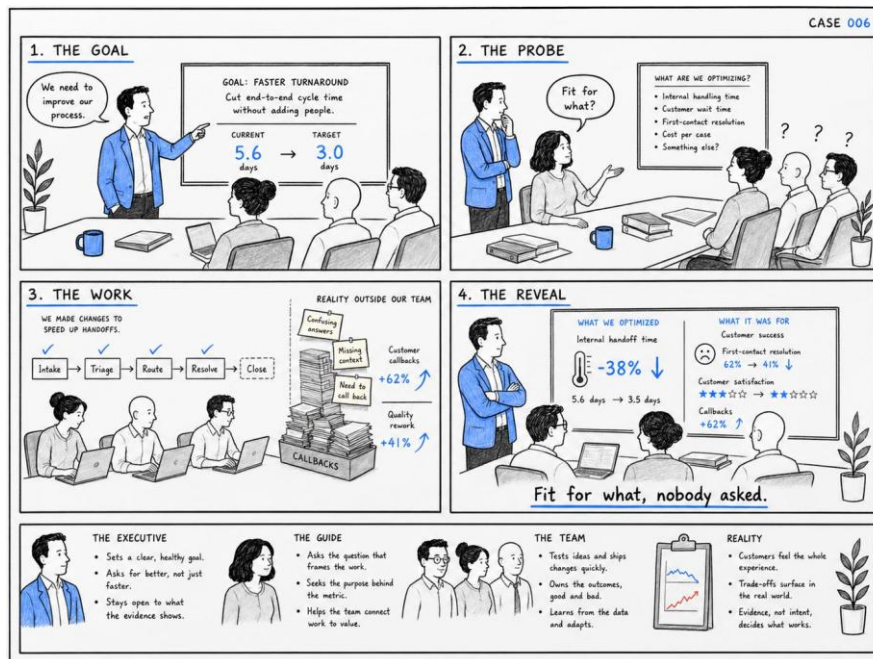
Part II — The Four Questions for Fitness

I once sat through a two-hour steering committee meeting in which twelve very senior people debated whether a target operating model was “fit for purpose.” Nobody, at any point, said what the purpose was. I don’t mean they disagreed on it. I mean it never came up. The phrase had become so smooth from overuse that it no longer needed a referent, the way “at the end of the day” no longer needs a day.

That meeting is the reason the Four Questions exist in this exact form and not some more elegant one. They are deliberately blunt, and they are the four things you have to answer before the word “fit” is allowed to mean anything: fit for what, by whom, how, and until when.

Fit for What

This is the one everyone thinks they’ve already answered. They haven’t. What they’ve usually answered is fit for the mandate — the thing written in the charter, the slide that got the budget approved. That is a real thing, but it is not the same as the actual condition the design has to hold up against, which is usually messier, more local, and less flattering than the mandate.




Case 006


FIELD NOTE — The phrase ‘fit for purpose’ is empty until the purpose is specific enough to exclude something. Ask what outcome, constraint, or stakeholder consequence the design must actually serve.

By Whom

This asks who is actually going to run this, live in it, be accountable when it wobbles — as opposed to who approved it. These are very often different people, and the gap between them is where most “change failed to stick” post-mortems quietly live.

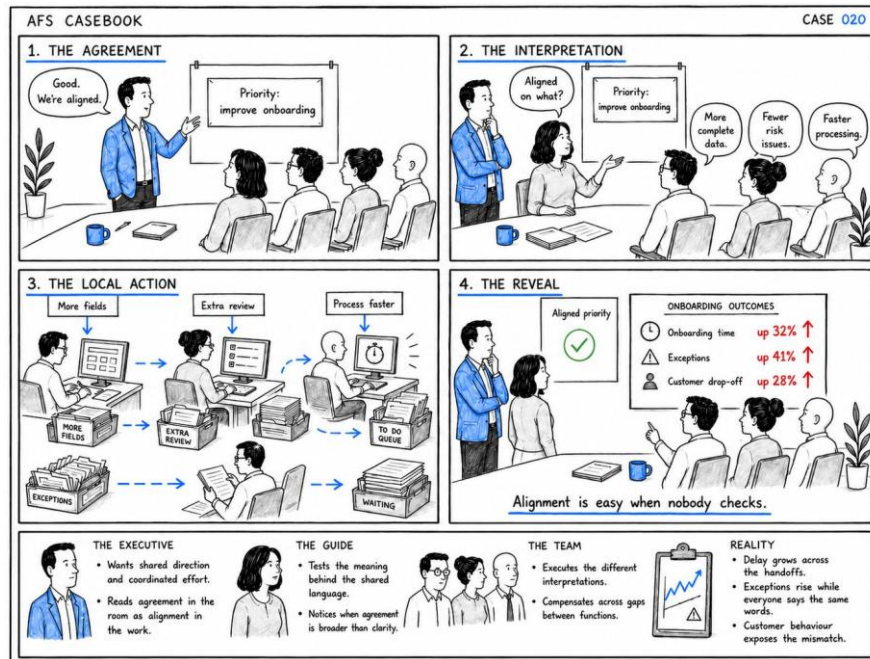


Case 026



Case 036

A related, cheerful little trap: everyone in the room agrees, out loud, and everyone in the room means something slightly different by the words they just agreed on. The agreement was real. The alignment was imaginary. Nobody checked, because checking would have required someone to ask a follow-up question, and follow-up questions have a way of ending meetings on time.

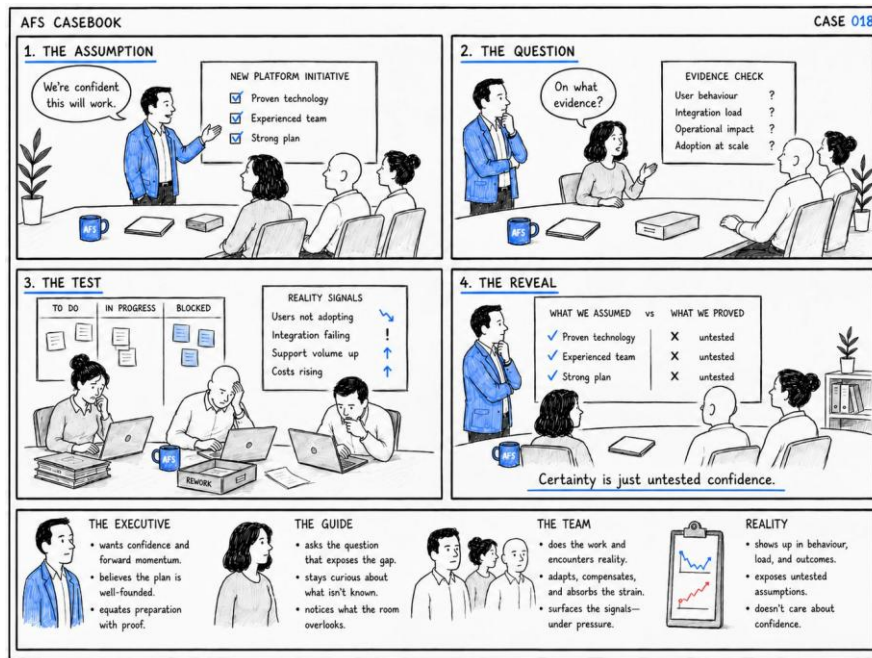


Case 020

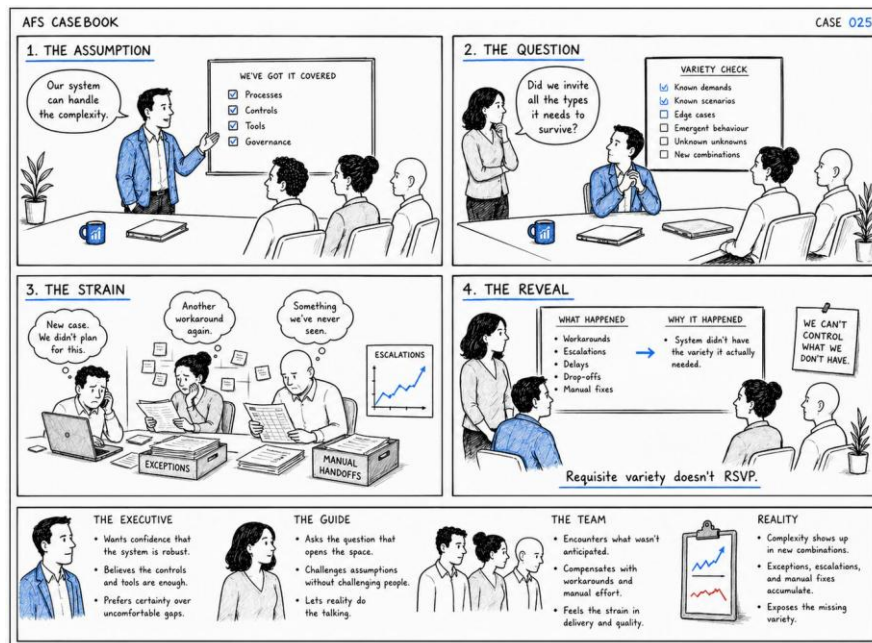
How

How — through what mechanism, what daily behavior, what actual sequence of actions — is the question organizations skip because it sounds operational rather than strategic, and strategic is where the prestige is. I've reviewed transformation decks with beautifully drawn target states and no credible account of how anyone gets from here to there on a Monday morning. Certainty, in these rooms, tends to arrive well before the evidence does, wearing a very confident expression.

And sometimes the honest answer to “how” is that the system needed a kind of variety — of skills, of scenarios, of edge cases — that nobody thought to invite. Complexity doesn't RSVP. It shows up whether or not there was a place set for it.



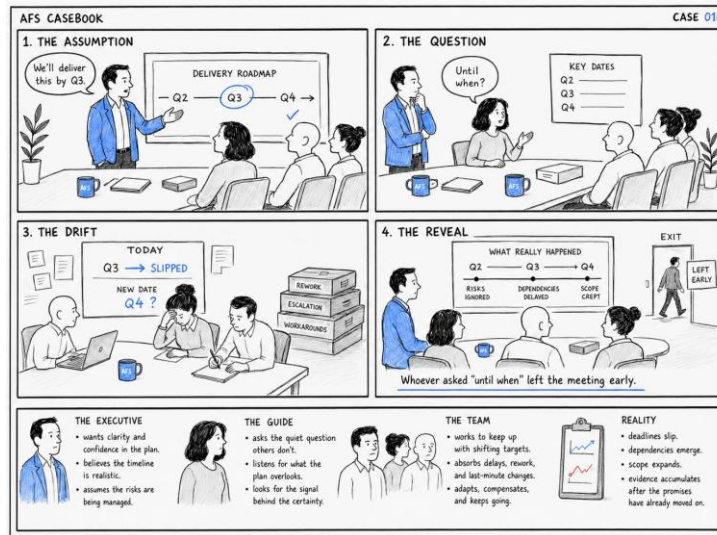
Case 018



Case 025

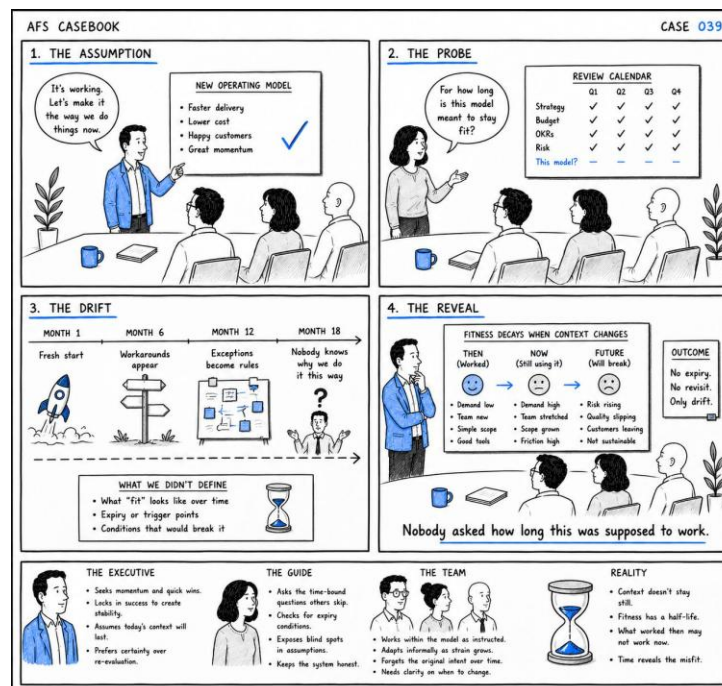
Until When

This is the question that admits, out loud, that fitness has an expiration date, and that pretending otherwise is the exact failure mode the Governing Law was built to catch. It is also, reliably, the first question that gets someone quietly checking their phone, because it's the one that implies future accountability rather than present-tense confidence.



Case 016

Sometimes nobody asks it at all, for eighteen months, at which point the organization discovers it has been running an operating model with no expiry date and no one who remembers agreeing to that.



Case 039

Here's the honest limitation, and I'd rather name it than have a sharp reader find it first: four questions are a discipline, not a guarantee. You can answer all four thoughtfully and still be wrong, because context moves faster than any assessment cycle. The questions don't eliminate self-deception. They just make it harder to hide, which is not nothing, but it isn't proof either.

Once you know what you're fit for, by whom, how, and until when, you need somewhere to put that information to work. That's Fitness-Driven Design.

Part III — From Questions to Design

There's a particular kind of org design review where everyone nods at the diagram and nobody asks how it was built. I don't mean the reporting lines — those get scrutinized to death. I mean the sequence of decisions that produced the diagram in the first place. Usually the honest answer is: someone had a model they liked, from a book or a prior company or a conference keynote, and the current organization got poured into it like batter into a pan shape that happened to be lying around.

Fitness-Driven Design is a practice, not another architectural layer. It means designing forward from the Four Questions instead of importing a structure because it worked elsewhere. The output is not 'the AFS design'. It is a context-specific design whose assumptions and failure conditions are visible.

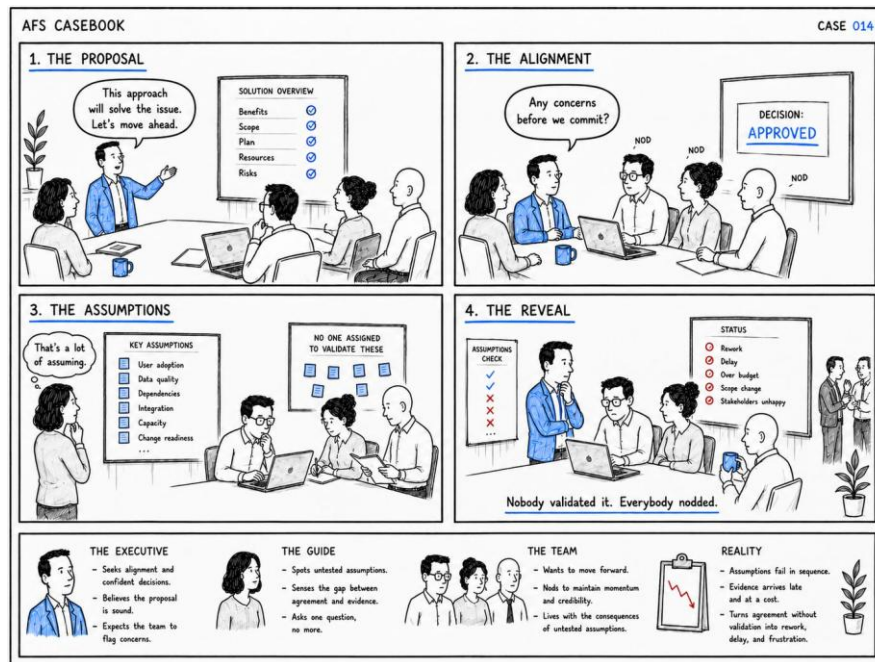
This sounds obvious stated flatly. It is rarely done, because designing forward from a diagnosis is slower and less reassuring than reaching for something with a name and a case study attached. "We did what worked at [respected competitor]" gets budget approved. "We built something specific to our own answers to four uncomfortable questions" gets you a lot of follow-up questions in the room.

Methods, unfortunately, do not produce evidence on their own. You can have every practice on the checklist ticked and still have nothing to show for it except a very tidy checklist.



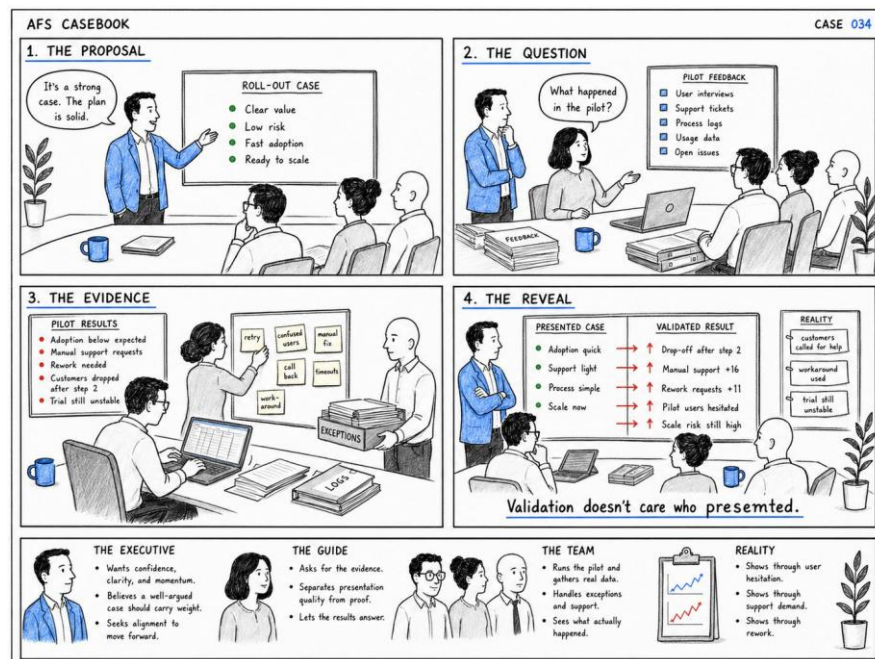
Case 004

And sometimes what actually happened is that a proposal moved through a room full of nodding heads, everyone agreeing to move ahead, and nobody — not one person — was assigned to go check whether the assumptions underneath it were true. Nodding, it turns out, is not validation. It only feels like it in the room.



Case 014


Reverse Validation is the practical safeguard: before building, state what evidence would show that the design has drifted out of fitness, then build that evidence-gathering into the work itself. It is not a separate ceremony. It is the discipline of making a design capable of contradicting its own story.



Case 034

It especially doesn't care whether the thing worked once. A pilot that goes well under a specific, generous set of conditions — extra staff, a priority freeze, someone's undivided attention — is a genuinely useful data

point. It is not, on its own, a method. “That worked once” and “that will always work” are different claims, and the gap between them is exactly where most scaled rollouts quietly fail.



Case 022

The useful outcome is system agility: the capacity of a design to notice when it no longer fits and to change without waiting for a ceremonial transformation. That capacity cannot be imported. It has to be designed into the way the system senses, decides, and acts.

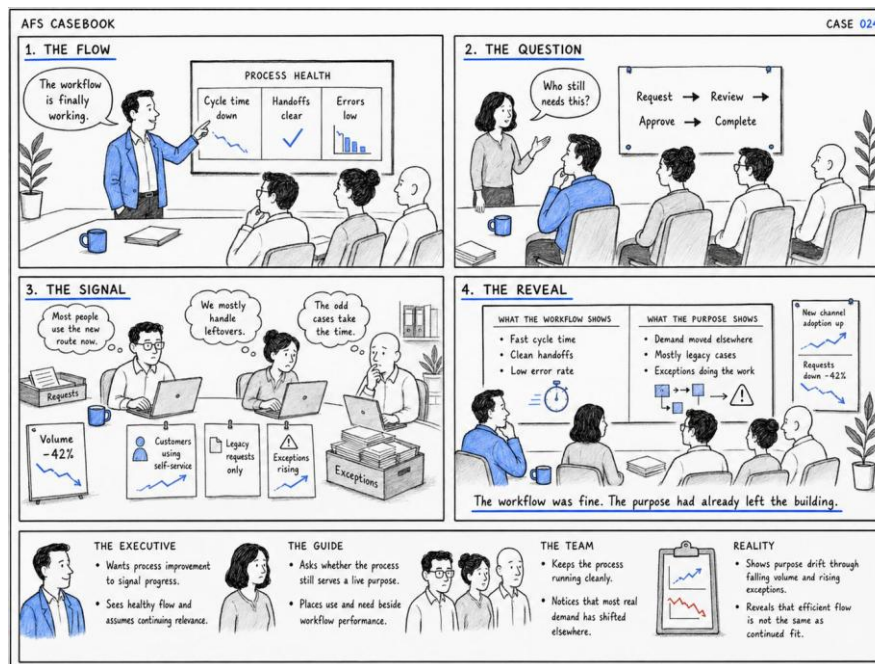
Part IV — The Fitness Sequence

Somewhere in most transformation programs there's a slide with four boxes and arrows between them, and everyone in the room already knows, without being told, that the arrows are aspirational. The boxes get built more or less in parallel under deadline pressure, and the arrows get drawn afterward to make the story coherent for the steering committee. I've drawn those arrows. I'm not proud of it. I understand exactly why it happens.

The Fitness Sequence is an argument against that habit. Fit for Purpose. Fit for Context. Fit for Execution. Fit for Improvement. Each one is the precondition the next one quietly assumes, and skipping a step is like skipping a rung on a ladder — you can do it, but you should know you're doing it.

Fit for Purpose

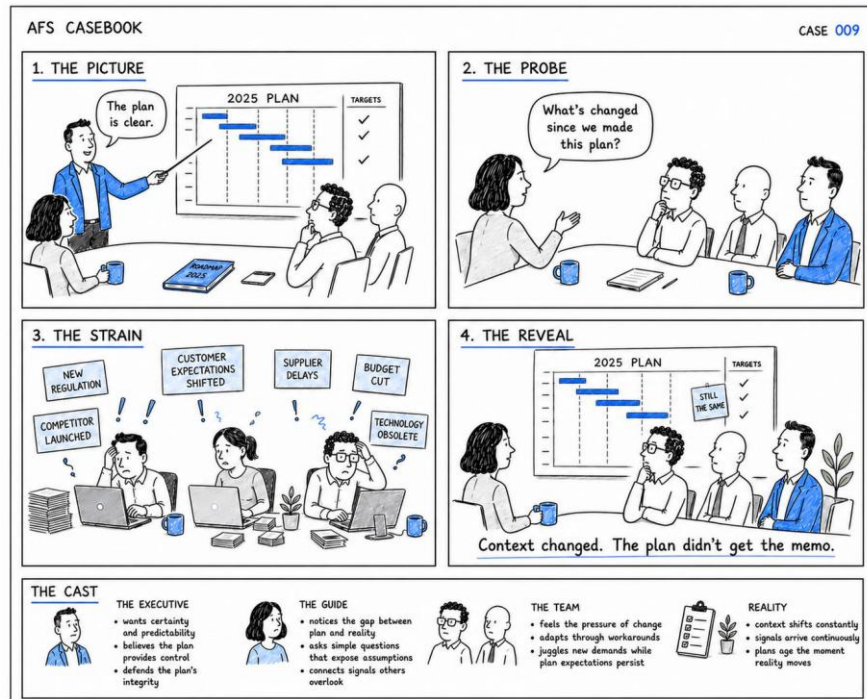
Does the design actually serve the What from the Four Questions — the real one, not the mandate's slogan version. This step feels complete on its own, which is exactly the trap. You can be entirely fit for purpose in the abstract and still discover the purpose itself has quietly relocated while nobody was watching the workflow, which kept humming along, clean and increasingly irrelevant.



Case 024

Fit for Context

This takes the purpose-fit design and asks whether it survives contact with the specific conditions it has to operate in — the actual people, actual history, actual politics, actual capacity. This is where most imported best practices die quietly, without anyone naming the cause of death. Context isn't scenery. It's load. And it changes at its own pace, entirely uninterested in whether the Gantt chart has caught up.



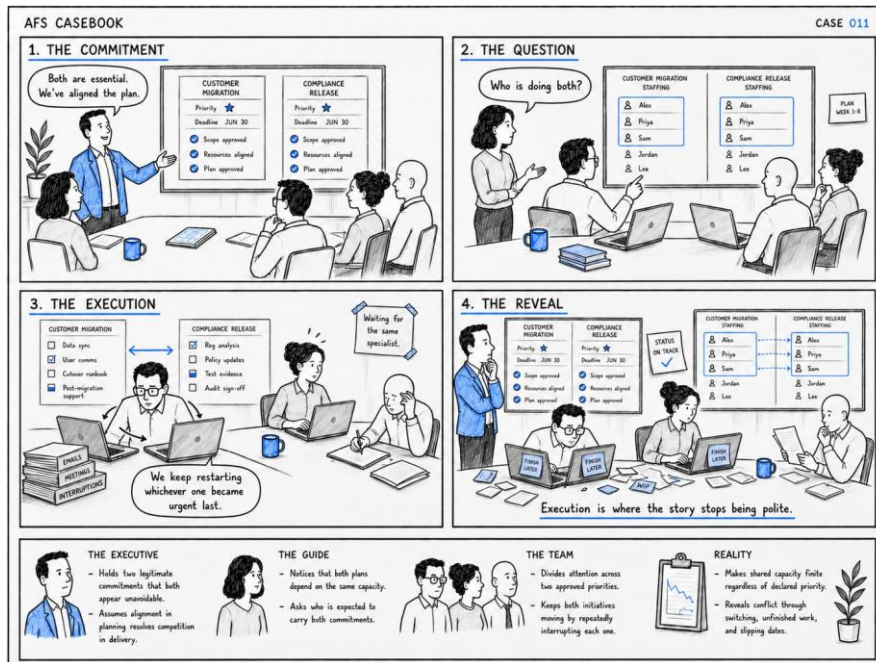
Case 009

FIELD NOTE — Context is not background information. It is part of the load the design must carry. List what has changed since the plan was approved and which assumption each change invalidates.


Fit for Execution

This is where the design meets Monday morning, and it's the step most decks skip entirely, because it's the least glamorous and the most mechanical. Can this actually be run, by the people who have to run it, with the skills, time, and authority they actually have — not the skills, time, and authority the design assumed they'd magically acquire once the new structure was announced.

A plan can be entirely fit for the meeting where it was presented — clear timeline, named owners, risks summarized, approvals in place — and entirely unfit for everything that happens after the meeting ends.



Case 011



Case 040

FIELD NOTE — Meeting fitness is not execution fitness. Ask what the plan requires on an ordinary Monday: time, skill, authority, coordination, and tolerance for exceptions.

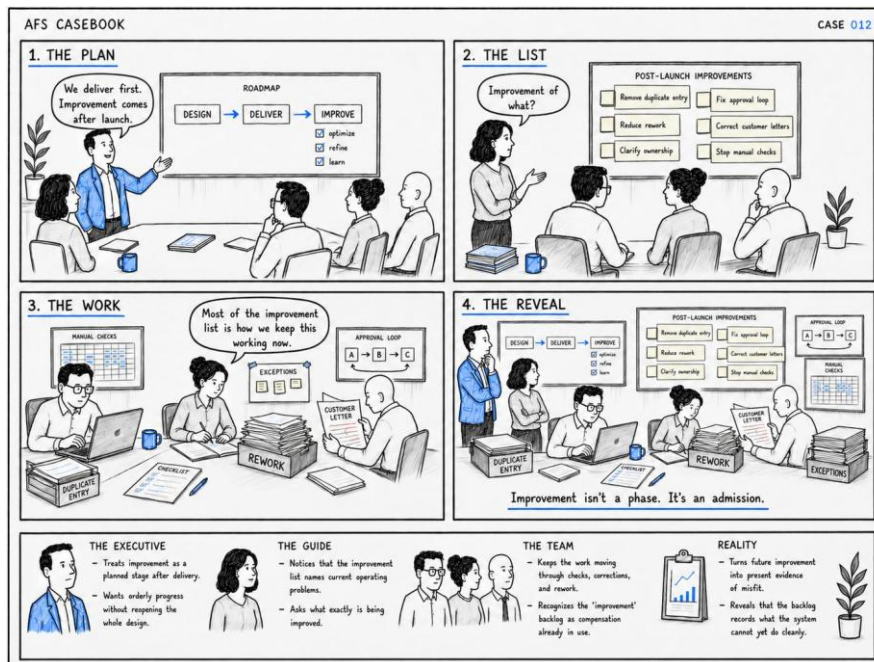
Fit for Improvement

This is the step that admits the first three were never going to be permanently correct, and builds in the mechanism for finding out when they stop being correct. It's also, not coincidentally, the step most likely to


be cut when budgets tighten, because it looks like it's checking something that's already “done.” Improvement isn't a phase that comes after delivery. It's a confession, in installments, about everything the design didn't quite get right the first time.

And its most honest form isn't a scheduled retrospective — it's the willingness to admit, at the exact moment certainty runs out, that this is where the actual learning starts.

The Sequence has an honest failure mode of its own. Organizations learn it well enough to perform it in order without doing the underlying work at each step. You can walk through all four stages in a workshop, tick every box, and still have produced nothing but a well-organized presumption. The stages don't move themselves. Something has to drive them, and that is the next part.



Case 012



Case 005

Part V — From Evidence to Action


I once asked a client what their strategy process actually was, mechanically, and after some throat-clearing the answer amounted to: someone senior gets a feeling, the feeling gets a name, and the name gets a budget line. I want to be fair — this is not stupidity. Most organizations don't have a mechanism for turning direction into motion. They have a moment of conviction, followed by a great deal of activity that hopes to be consistent with it.

Direction becomes adaptation through a loop with three stages. The Adaptive Signal Field is where the organisation reads what has changed and how much it matters. The Friction-Fitness Map locates that reading: where the pressure sits, and whether it is telling the truth. The AFS Adaptive Engine is the organisation's actual ability to turn a located reading into coordinated action. The loop matters more than its labels.


Read what has changed. Locate where it presses. Act through the real system of work. Then examine what the action revealed, because that is the next signal.

The Adaptive Signal Field

The Adaptive Signal Field orients: given the purpose, context, evidence, and constraints, which direction now has the strongest claim? A reading is not a decision. It is an explicit directional hypothesis that can still be ignored, distorted, or overtaken by events. The cartoons that follow still use older words for this, a compass for direction and a dial for scale. The words retired. The failure modes did not.




Case 030



Case 037

The same field calibrates: how much correction, applied how quickly, with what reversibility? Some signals require a small experiment; others expose a design that can no longer be defended. Treating every signal as either noise or emergency is a failure of calibration.




Case 019

The Friction-Fitness Map


Between the reading and the action sits the question of where. A signal nobody can locate gets addressed everywhere, which in practice means nowhere. The Map gets its own part later in this book. Here it is enough to know that the loop passes through it, and that skipping it is how organisations end up fixing the loudest team instead of the misfit one.

The AFS Adaptive Engine

Action is where the loop either becomes real or collapses into diagnosis theatre. The organisation must be able to coordinate across boundaries, allocate authority, clear blockers, and alter the work. That capability often exists informally before the org chart recognises it.



Case 017



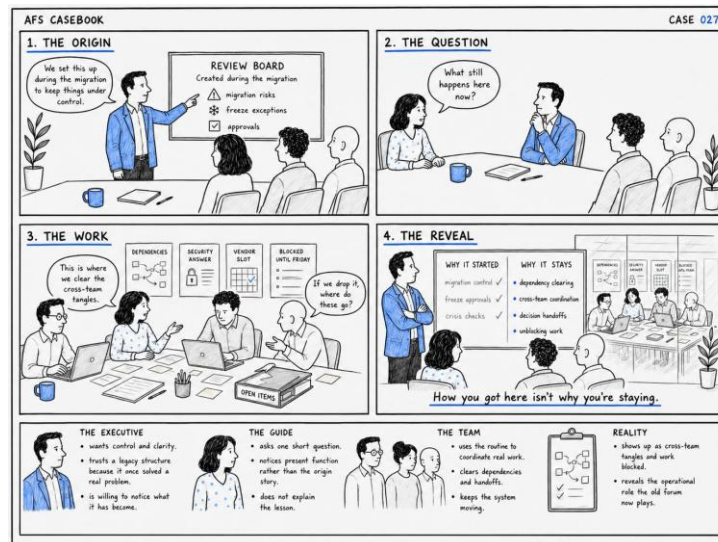
Case 031

Sometimes the adjustment is smaller and less visible than a whole engine turning over — just a quiet, unofficial workaround, adopted because it kept the work moving, that nobody with the authority to approve it ever actually saw.



Case 042

And sometimes a mechanism outlives the reason it was built entirely, and keeps running for reasons that have nothing to do with its origin story and everything to do with the fact that people are still, genuinely, using it to get things done.



Case 027

A brilliant reading connected to no capacity for action is an accurate prophecy. The FVI Lens, which the next part takes up, is not a fourth stage in this loop. It is the lens the whole loop runs through, and it keeps the reading rigorous rather than reducing it to 'look at the dashboard and feel a mood'.

The 20-Minute Fitness Check

Use this when a room is about to approve, continue, scale, or defend something described as ‘working’. The point is not consensus. The point is to expose the claim and the evidence quickly enough to affect the decision.

Minutes	Question	Output
0–3	What exactly are we claiming is fit?	One sentence naming the design, purpose, and current conditions.
3–7	Fit for what, by whom, how, and until when?	Four answers specific enough to be challenged.
7–12	What evidence supports the claim? What evidence would change our mind?	Supporting evidence plus at least one disconfirming indicator.
12–16	Where is friction visible, hidden, or being mistaken for healthy tension?	One pressure point that deserves inspection.
16–20	What will we check, change, or stop—and when?	One owner, one next check, one expiry or review condition.

Failure condition: if the room cannot name evidence that would change its mind, it is not validating the claim. It is defending it.


Part VI — The FVI Lens

There's a specific look a room gives you when you ask “how do you know this is working,” and it's not defensiveness, exactly. It's more like the look you get asking someone how their marriage is going at a dinner party — a flicker of “that's not really a question with an answer, is it,” followed by something warm and vague about how things feel pretty good these days. Organizations answer “how do you know” with sentiment dressed as measurement, and they do it fluently enough that nobody in the room notices the substitution happening.

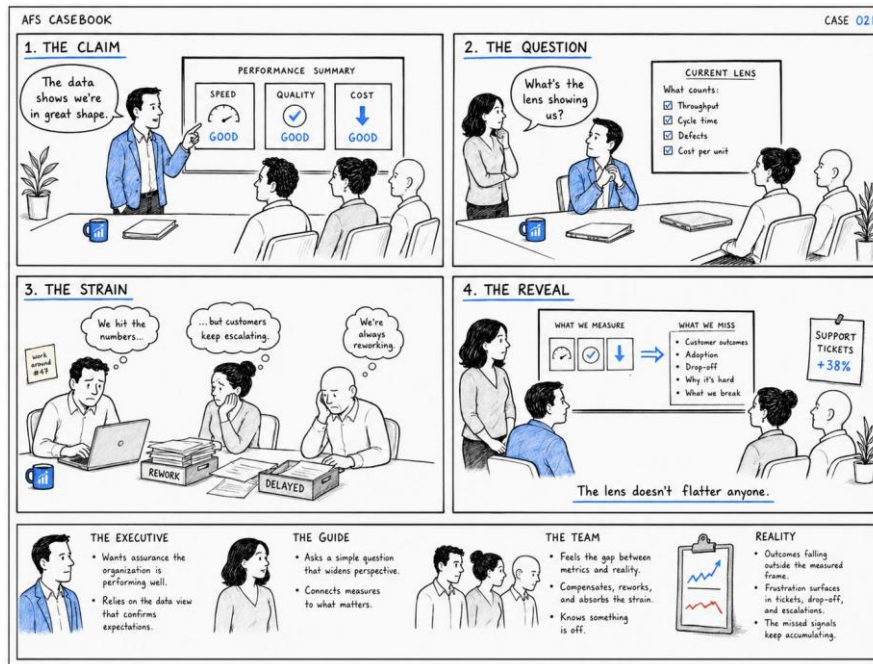
The FVI Lens separates three questions that organisations routinely collapse into one impression: does the design currently fit; is that claim supported by evidence; and what should that evidence change about the next decision? Fitness, Validation, and Intelligence can disagree. That disagreement is information, not a reporting defect.

Fitness

This aperture asks the question the whole approach is named for: given current conditions, does this still hold up. Dashboards are very good at showing you activity and rather bad at showing you fit. A dashboard can report every indicator green while the support queue quietly climbs into the hundreds one desk over — and, to be fair to the dashboard, it wasn't lying. It just wasn't asked the right question.



Case 002

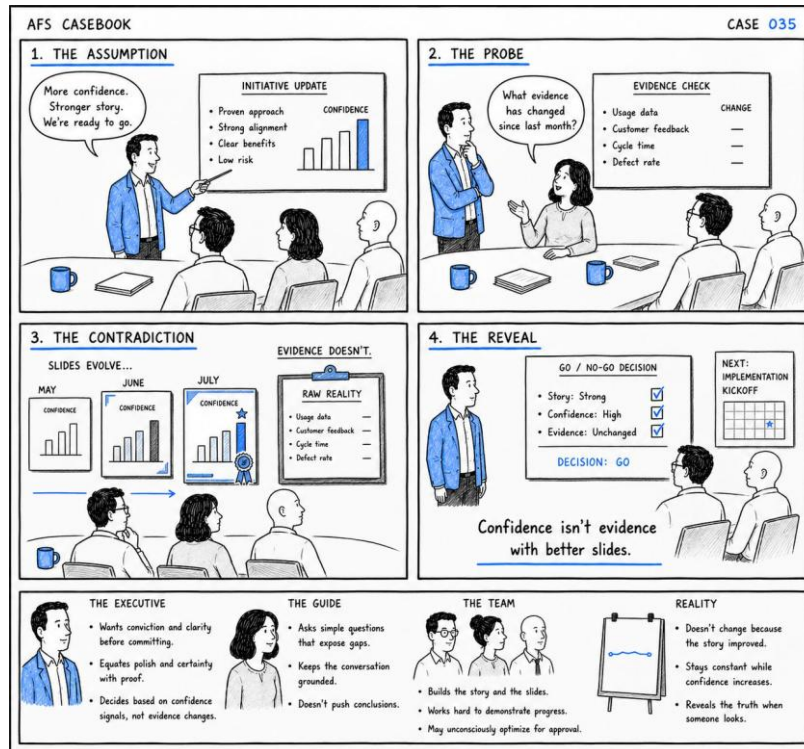


Case 021

Validation

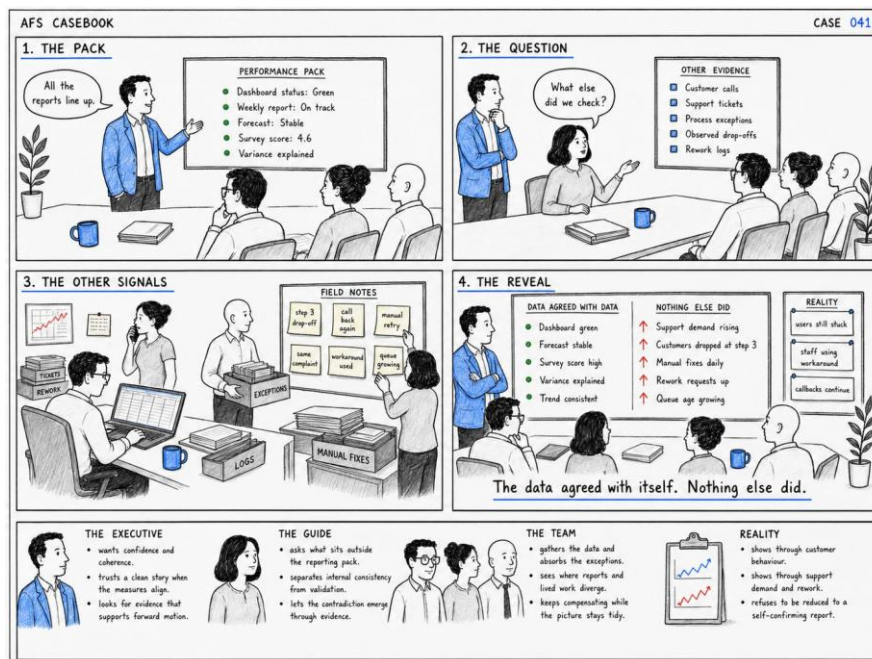
This asks something adjacent but distinct: is there actual evidence for the claim of fitness, or is it a claim resting on authority, tenure, or the fact that nobody's complained loudly enough yet. Confidence is not evidence, however good the slides look, and however much the story has visibly improved between the May version and the July one with the little gold star on it.

Sometimes the data is internally spotless — every reported figure agrees with every other reported figure — and every piece of evidence sitting outside that particular report is climbing steadily in the wrong direction, unconsulted.



Case 035

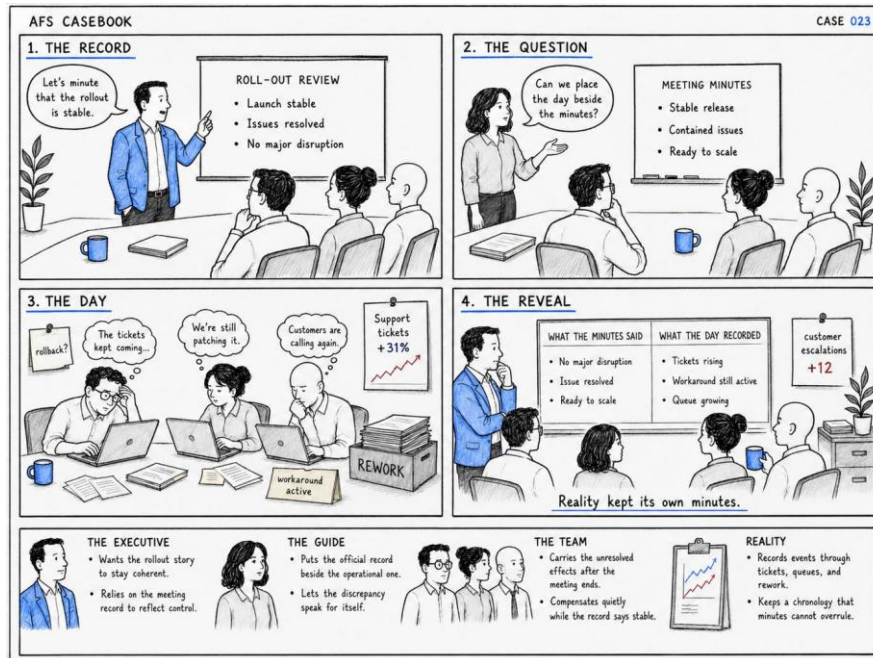
FIELD NOTE — Confidence is a property of the presenter. Validation is a property of the claim. Ask what evidence could survive a change of presenter.



Case 041

FIELD NOTE — Internal consistency is not external validity. Identify one operational source outside the reporting chain and check whether it tells the same story.

And sometimes the official record and the operational one simply part ways, politely, in the same meeting. The minutes say stable. The support queue keeps its own minutes, and doesn't submit them for approval.




Case 023


Intelligence

This asks the forward-looking question: given what Fitness and Validation are showing you, what does this predict, and what should you do differently because of it. This is the aperture most often skipped, because Fitness and Validation are diagnostic — comfortable, even when the diagnosis is bad — and Intelligence demands you commit to an implication that someone might later check you against.

Sometimes intelligence arrives well ahead of the official meeting that's supposed to discuss it, quietly, in a signal that keeps recurring before anyone's given it a name.




Case 044



Case 003

And sometimes it just sits there, unopened, in an inbox nobody's checking, while the thing it was warning about keeps happening anyway, on schedule, whether or not the door ever gets opened.



Case 029

The honest limitation here: the FVI Lens requires a kind of organizational patience that most reporting cadences actively punish. Boards want a status. Green, amber, red. The Lens wants to tell you that Fitness is amber, Validation is green, and Intelligence hasn't been asked yet — which is a true and useful answer that does not fit in a single traffic-light cell. Looking through the Lens tells you where things stand. It doesn't yet tell you where the pressure is building. That's what the Friction-Fitness Map is for.

The Fitness Claim Record

A fitness conversation becomes useful when its assumptions survive beyond the meeting. Capture the minimum below; do not turn it into a compliance form.

Field	Record
The claim	What, exactly, is currently believed to be fit?
Conditions	For what purpose and context is the claim supposed to hold?
Operators	By whom is it actually executed, maintained, and experienced?
Mechanism	How does it work in ordinary practice—not in the presentation?
Evidence	What supports the claim now?
Disconfirmation	What observation would make us revise or withdraw the claim?
Expiry	Until when is the claim credible, or what change triggers review?
Next action	What will be checked, changed, stopped, or tested next—and by whom?

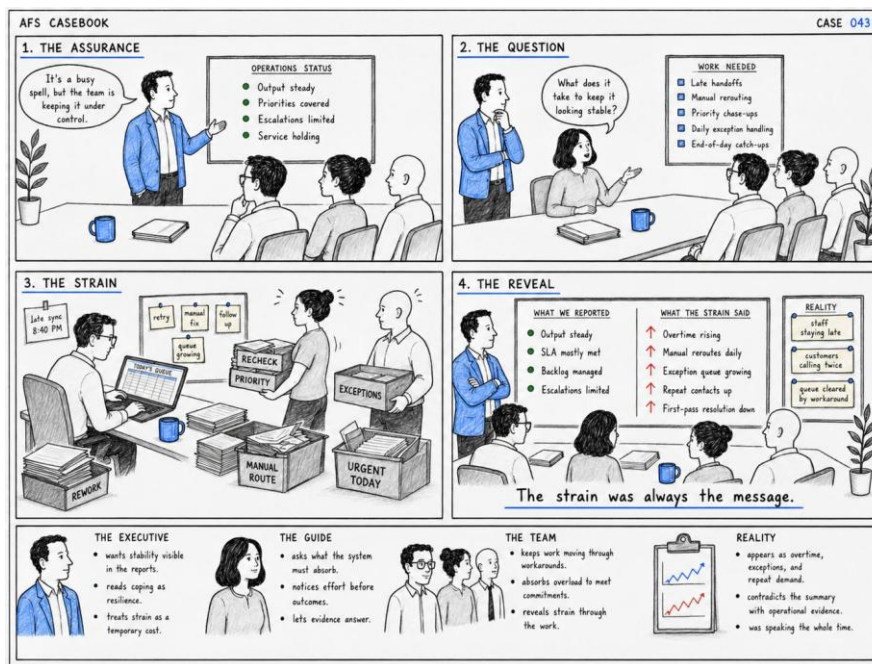
Part VII — The Friction-Fitness Map

Every organization has at least one team that everyone quietly agrees is “just difficult,” in the tone people use for a relative whose politics nobody wants to discuss at dinner. Nobody investigates why. It's simpler to treat friction as a personality trait of the team than as a signal about the system they're operating inside.

The Friction-Fitness Map exists to stop that reflex. It places friction against two axes — how much strain is present, and whether that strain is visible or hidden — producing four quadrants: System Strain, Adaptive Tension, Hidden Misfit, and Stable Flow.

System Strain

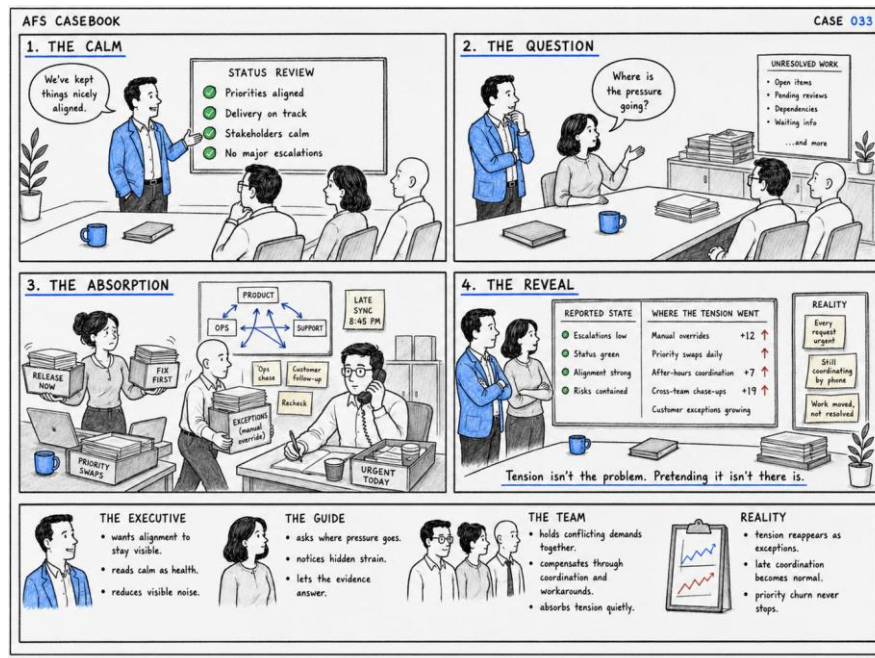
This is friction that's loud and structural — everyone can feel it, and it traces back to the design itself. It's the easiest quadrant to work with, uncomfortably, because at least nobody's pretending otherwise. The strain tends to speak in a very consistent language: overtime, escalations, repeat contacts, the kind of evidence that was never really trying to be quiet.



Case 043

Adaptive Tension

This is friction that's loud but not structural — the discomfort of a system that's actually adapting well, which is rarely comfortable. This quadrant is most likely to be mistaken for System Strain and “fixed” prematurely, which is its own kind of harm. Sometimes tension isn't the problem. Pretending it isn't there is.



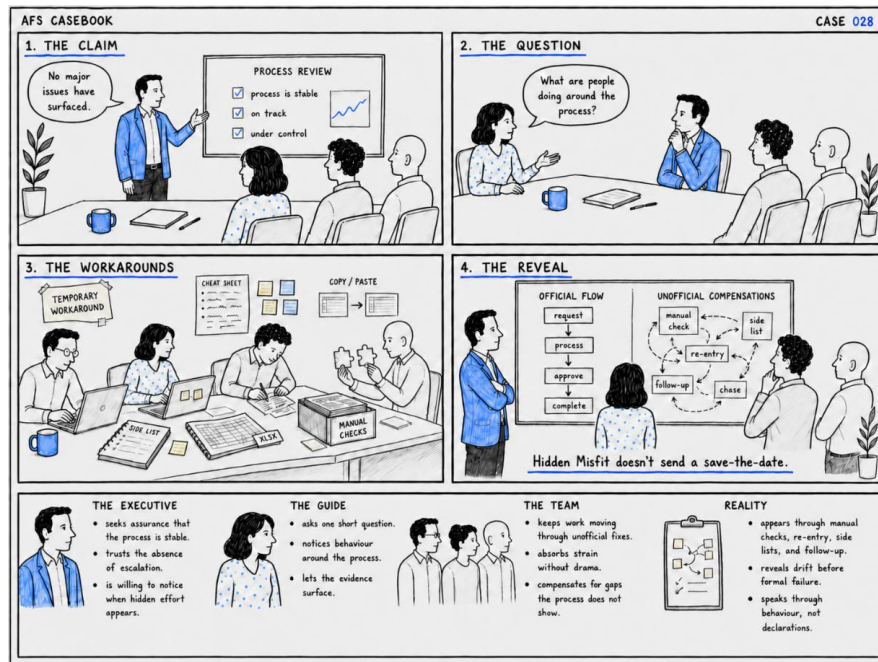
Case 033

FIELD NOTE — Do not remove all tension. First ask whether the tension is exposing misfit or producing adaptation. The same discomfort can require opposite interventions.

Hidden Misfit


This is the dangerous one — low visible strain, but the fit is actually gone. Everything looks calm because the people affected have stopped signalling, usually because signalling stopped producing any response worth the effort. It doesn't send a save-the-date. It just quietly accumulates behind a process that looks, from the front, entirely stable.

An area can look genuinely broken and, at the same time, not actually be failing at anything, because nothing was ever asked of it that could fail. Nothing here is broken. Nothing here works, either. Both statements can be true at once, and the room usually only clocks the first one.




Case 028

FIELD NOTE — Silence is ambiguous. It can mean flow, resignation, invisibility, or successful workaround. Ask what happened the last two times people raised the strain.



Case 010

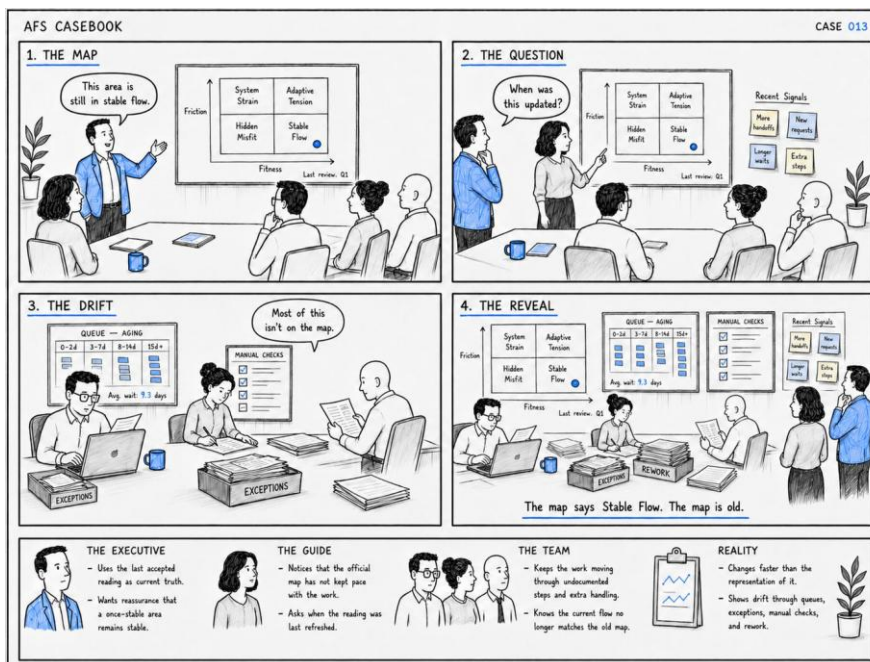
And a structure can hold this quadrant for a very long time before anyone notices, mostly because the org chart it's drawn on has no mechanism for noticing anything.



Case 007


Stable Flow

This is the quadrant everyone wants to believe they're in, and the one the Governing Law should make you suspicious of by default — not because Stable Flow doesn't exist, but because claiming it without evidence is exactly the presumption AFS was built to catch. Real Stable Flow has been checked. Claimed Stable Flow that hasn't been checked is usually Hidden Misfit wearing a nicer outfit, standing in front of a map that was accurate the last time anyone updated it.



Case 013

Sometimes the map has simply gone stale, and the friction underneath it has been dutifully filing its own reports the whole time, in the form of exceptions, manual checks, and everything the official process quietly stopped asking about.



Case 015

A short illustration, from a banking context I worked in some years back, without naming names or getting more specific than the point requires. A regional operations team had, by every visible measure, gone quiet — few escalations, stable throughput, no complaints reaching leadership. Read at face value, that's Stable Flow. Read through the FVI Lens, the Validation aperture told a different story: nobody could point to evidence for why things were calm, only the absence of noise. A short round of structured conversations placed the team firmly in Hidden Misfit — a workaround from a prior system migration had become permanent, unofficial, and increasingly brittle, and the team had simply stopped reporting the strain because two previous attempts to raise it had gone nowhere. The friction hadn't disappeared. It had gone underground, which is worse.

The honest limitation of the Map: quadrants are diagnostic categories, not fixed addresses. A team can be in Adaptive Tension on one dimension of their work and Hidden Misfit on another, simultaneously, and the Map doesn't stop you from oversimplifying into a single quadrant label the way it warns you against oversimplifying fitness into a single RAG status.

That's the operating architecture of AFS, start to finish — the Governing Law down through the instruments that keep it honest in practice. What's left is the harder question, and the one this whole book owes you by its own logic: does AFS hold up to its own law. That's the close.

Lineage, Boundaries, and Debt

AFS did not emerge from an intellectual vacuum. It sits in a lineage that includes systems thinking, cybernetics, organisational learning, requisite variety, sensemaking, evidence-based management, dynamic capabilities, and Fit for Purpose. Its contribution is not the invention of adaptation, validation, or fit. It is the configuration of these concerns around one governing discipline: a claim of organisational fitness must remain open to disconfirmation.

That lineage also imposes boundaries. AFS is not a substitute for domain expertise, strategy, ethics, technical design, or political judgement. It does not tell an organisation what purpose to choose. It cannot make weak evidence strong, remove conflicts of interest, or guarantee that people with authority will act on what they learn.

Its current evidence base is primarily practitioner experience and repeated pattern recognition. That is enough to justify inquiry and application; it is not enough to claim universal validity or causal proof. The next stage of the work must include documented applications, counterexamples, comparison with ordinary review practices, and clearer boundary conditions.

What Would Prove AFS Unfit?

AFS should be revised, narrowed, or abandoned where credible evidence shows any of the following:

- It produces no better questions or decisions than an ordinary well-run review.
- Practitioners cannot distinguish its core concepts reliably when examining the same situation.
- The Four Questions add ceremony but do not change evidence, action, or accountability.
- The vocabulary encourages delay, abstraction, or diagnostic theatre rather than intervention.
- Organisations repeatedly use AFS to legitimise decisions already made.
- Its diagnoses do not transfer beyond the settings in which the approach was developed.
- Its use creates more interpretive complexity than practical clarity.

These are not rhetorical risks. They are a provisional test programme for the approach itself.

Closing — The AFS Paradox

Somewhere around the third or fourth time I presented this material, a participant asked a question I'd been quietly hoping nobody would ask, mostly because I didn't yet have a good answer. She said: you've built an approach whose entire premise is that fitness must be proven and not presumed. How do you know AFS is fit? Isn't that exactly the claim you're not allowed to make about yourself without proving it?

She was right, and the honest answer took me longer to arrive at than I'd like to admit. The comfortable move would have been to treat AFS as the instrument and everything else as the thing being measured — the Lens doesn't need its own lens, the argument goes, the same way a ruler doesn't need to measure itself. But that argument only works if AFS is a neutral measuring tool, and it isn't. It's a designed thing, built by a person, in specific contexts, with specific blind spots, and it makes the same claim on its own behalf that it makes on everyone else's: that fit is not a status you earn once, it's a condition you keep having to demonstrate. There is no exemption clause in the Governing Law that reads “except for the law itself.”

So here is the paradox stated plainly, because stating it plainly is the entire point: an approach built to detect presumed fitness must itself avoid presuming its own fitness, and it cannot simply declare that it has passed this test, because declaring it would be the exact failure mode it exists to prevent. AFS has to be willing, continuously, to ask of itself the same four questions it asks of everything else. Fit for what, exactly, and has that “what” drifted since the approach was built. By whom is it actually usable, and is that still the same population it was designed for. How is it actually practiced, in the room, versus how it's described in a book like this one. Until when does this version of AFS remain the right one, and what would tell us it's expired.

I do not have final answers to those questions. What I have is a practice of asking them, a growing body of cases, and a responsibility to document where the approach changes decisions—and where it does not. Reciting the Governing Law is not evidence that the Law is being practised. I have seen that failure in others and in myself.

That is not a triumphant ending. It is a more useful one. AFS remains a practitioner hypothesis with a disciplined way of examining itself. Its future value will depend less on how coherent it looks in this book than on whether people can use it to expose a weak claim, improve a decision, and notice when the approach itself has become part of the theatre.

The Executive would like a tidier finish. The Guide would ask one more short question and let the silence do the rest. The Team is already back at their desks, quietly keeping the whole thing running regardless of how this chapter ends. And Reality — Reality was never going to RSVP to a closing chapter either. It'll be back Monday, same as always, with its own minutes.

Fitness is proven, not presumed — including this sentence, including this book, including whatever certainty you're feeling about all of it right now, on this page, which will need checking again before too long.

Casebook Index

All forty-five cases in numerical order. The “Illustrates” column names the practical concept or question each case is best suited to examine.

Case	Illustrates	Caption
000	Governing Law	Fitness is proven, not presumed.
001	Governing Law	Reality doesn't attend announcement meetings.
002	FVI Lens — Fitness	The dashboard wasn't lying.
003	FVI Lens — Intelligence	The signal arrived first.
004	From Questions to Design	Methods don't produce evidence.
005	Fitness Sequence — Improvement	Improvement starts where certainty ends.
006	Four Questions — What	Fit for what, nobody asked.
007	Friction-Fitness Map — Hidden Misfit	The org chart doesn't know it's wrong yet.
008	Governing Law	Presumed fitness has excellent posture.
009	Fitness Sequence — Context	Context changed. The plan didn't get the memo.
010	Friction-Fitness Map — Hidden Misfit	Nothing here is broken. Nothing here works, either.
011	Fitness Sequence — Execution	Execution is where the story stops being polite.
012	Fitness Sequence — Improvement	Improvement isn't a phase. It's an admission.
013	Friction-Fitness Map — Stable Flow	The map says Stable Flow. The map is old.
014	From Questions to Design	Nobody validated it. Everybody nodded.
015	Friction-Fitness Map — System Strain	The friction was the feedback.
016	Four Questions — Until When	Whoever asked “until when” left the meeting early.
017	Evidence to Action	The system adapted. The org chart didn't notice.
018	Four Questions — How	Certainty is just untested confidence.
019	Evidence to Action — Adaptive Signal Field	The dial moved. Nobody was watching the dial.
020	Four Questions — By Whom / How	Alignment is easy when nobody checks.
021	FVI Lens — Fitness	The lens doesn't flatter anyone.
022	From Questions to Design	That worked once. Once isn't a method.
023	FVI Lens — Validation	Reality kept its own minutes.
024	Fitness Sequence — Purpose	The workflow was fine. The purpose had already left the building.
025	Four Questions — How	Requisite variety doesn't RSVP.
026	Four Questions — By Whom	Fit for whom wasn't on the agenda.
027	Evidence to Action	How you got here isn't why you're staying.
028	Friction-Fitness Map — Hidden Misfit	Hidden Misfit doesn't send a save-the-date.
029	FVI Lens — Intelligence	Intelligence arrived. Nobody opened the door.
030	Evidence to Action — Adaptive Signal Field	The compass pointed somewhere. The room kept walking the old way.

Case	Illustrates	Caption
031	Evidence to Action — Action	The engine was running. The org chart assumed it was idling.
032	Governing Law	Everyone's working harder. Nothing's getting fitter.
033	Friction-Fitness Map — Adaptive Tension	Tension isn't the problem. Pretending it isn't there is.
034	From Questions to Design	Validation doesn't care who presented.
035	FVI Lens — Validation	Confidence isn't evidence with better slides.
036	Four Questions — By Whom	By whom was assumed. By whom was wrong.
037	Evidence to Action — Adaptive Signal Field	The compass was right. The org kept consulting the old map.
038	Governing Law	The org asked reality to wait its turn.
039	Four Questions — Until When	Nobody asked how long this was supposed to work.
040	Fitness Sequence — Execution	The plan was fit for the meeting. Nothing else.
041	FVI Lens — Validation	The data agreed with itself. Nothing else did.
042	Evidence to Action — Action	Something's adjusting. Nobody signed off on it.
043	Friction-Fitness Map — System Strain	The strain was always the message.
044	FVI Lens — Intelligence	Intelligence without validation is just an opinion with a chart.