The Diagnosis Was Correct

Two Transformations. One Honest Retrospective.

By Mario Aiello – Adaptive Ways

There are two ways to be wrong about an organizational transformation. The first is to misread the situation, get the diagnosis wrong, misidentify the constraints, apply the wrong instruments. The second is to read the situation correctly and watch it fail anyway.

This paper is about the second kind.

The two engagements described here took place in different institutions, a major bank and an insurance company, separated by context, culture, and the specific shape of their failure. What they share is this: a correct diagnostic reading of the organizational situation, a set of interventions that produced genuine results at the team level, and a macro transformation that did not survive contact with the organizations that commissioned it.

The Adaptive Fitness System did not exist as a named approach when these engagements happened. But the diagnostic logic did, the Four Questions, the attention to context over method, the insistence that fitness is always relative to demand. Applying that logic retrospectively to both cases is an honest exercise, not a flattering one. The question is not whether the AFS would have diagnosed these situations correctly. It would have. The question is what correct diagnosis is actually worth when the organization isn't structurally ready to receive it.

The answer, in both cases, is less than you might hope.

Case One: The Bank

The bank had scale, sponsorship, and ambition. Senior leadership had decided that transformation was necessary, that the organization needed to become more adaptive, more responsive, more capable of delivering value at speed. These were not unreasonable conclusions. The operating environment demanded them.

What leadership did not do was give itself the means to achieve what it had decided to want.

The means it chose was a copied scaling framework, one of the large, branded models that had been successfully deployed, or at least successfully marketed, in organizations that resembled this bank in size if not in character. The framework arrived with its ceremonies intact, its roles defined, its cadence pre-specified. It was applied to a hierarchical culture, a matrix ownership structure, and a legacy governance apparatus that had not been touched. The framework was deployed into an organization that had not been redesigned to host it.

The result was predictable in retrospect and was not predicted at the time. Agility became local. Teams adopted the vocabulary, ran the ceremonies, produced the artefacts. At the system level, the command-and-control reflex remained intact. Leadership remained disconnected from delivery. Agendas remained split between those who wanted transformation and those who wanted the appearance of transformation. The framework spread; fitness did not.

What the AFS Would Have Read

The Four Questions, applied at the outset, would have produced an uncomfortable reading.

Fitness for Purpose was rhetorical rather than committed. The stated intent, genuine adaptive capability, was real at the level of aspiration. But the means chosen revealed the actual intention, which was to demonstrate transformation without surrendering the organizational conditions that made transformation impossible. A leadership team that wants adaptive capability but will not redesign its governance structure does not actually want adaptive capability. It wants something that resembles it from a sufficient distance.

Fitness for Context was low. The scaling framework had been designed for a different kind of organization, one with sufficient autonomy at the team level, sufficient alignment at the portfolio level, and sufficient trust between leadership and delivery to make the framework's assumptions hold. None of those conditions were present. The framework travelled, but its prerequisites did not.

Fitness for Execution was uneven. Some teams had real coaching, real product ownership, real flow. Others had the ceremonies without the substance. The unevenness was not an implementation failure. It was the expected output of an approach applied without contextual adaptation, some contexts absorbed it, others rejected it, and the framework had no mechanism for distinguishing between the two.

Fitness for Improvement was near zero at the system level. Feedback loops existed within teams. Between teams and leadership, almost nothing moved. The learning that happened locally did not propagate upward. The organization could not sense what was working and what was not, because the sensing mechanisms had not been built.

This diagnosis would have been available at the start of the engagement. It was available, in a less structured form, to anyone paying attention. The question is what it was worth.

The Limit of Early Diagnosis

If the AFS diagnostic had been applied formally at the outset, if the fitness reading had been delivered to leadership before the scaling framework was committed to, there are two plausible outcomes. One is that the diagnosis lands and changes the direction: the framework is abandoned or substantially modified, the governance structure is addressed, the conditions for transformation are built before the transformation is attempted. The other is that the diagnosis

is received politely, noted, and set aside, because the decision to use the framework had already been made by people whose authority derived from making exactly this kind of decision, and who had chosen it for reasons that a fitness reading could not address.

The honest answer is that the second outcome was more likely. Not because the diagnosis would have been wrong, but because the people who needed to act on it were the same people who had already decided what they wanted the answer to be. A correct reading of organizational unfitness does not dissolve the organizational dynamics that produced the unfitness. It names them. What happens next depends on whether anyone with the authority to act on the naming is willing to do so.

In this case, they were not. The transformation proceeded as planned. The framework was deployed. The micro results were real. The macro ambition remained an illusion.

Case Two: The Insurance Company

The insurance company was a different kind of situation. The bank's failure was an illusion, leadership had the dream but not the means, and the gap between aspiration and commitment was wide enough to drive a transformation through and lose it. The insurance company was something harder: a premise that was structurally impossible from the start.

The premise was this: agile is something you adopt without changing how the organization works.

This is not an uncommon premise. It is, in fact, the premise on which most failed agile adoptions rest. But in this case it was unusually explicit. The organization had decided, before any engagement began, what agile meant and what it did not mean. It meant teams running sprints, producing backlogs, attending stand-ups. It did not mean touching governance. It did not mean redesigning accountability structures. It did not mean asking leadership to change its relationship to delivery. The perimeter of the intervention had been drawn in advance, and the perimeter excluded everything that would have made the intervention consequential.

The pilot was introduced, not invited. Teams participated because they had been directed to, not because they had chosen to. Management measured the output of the pilot using the same control-oriented metrics they used to measure everything else. Feedback loops that existed within teams were cut off at the reporting layer. Budget decisions terminated the learning cycles that had begun to produce results.

This was not resistance to change in the ordinary sense. It was something more structural: an organization that had already decided what agile was, and whose decision made genuine agile adoption definitionally impossible. The initiative could produce compliance. It could not produce fitness.

What the AFS Would Have Read

The fitness reading here would have been brief and unsparing.

Fitness for Purpose was absent. The agile initiative was framed as a delivery improvement, a way of producing the same outputs more efficiently. It was not framed as a capability development effort, and it was not understood as one. The purpose was too narrow to support the intervention being attempted.

Fitness for Context was structurally blocked. The organizational context, hierarchical culture, command governance, compliance-oriented management, was not an obstacle that could be worked around. It was the operating system. An adaptive intervention requires a context that can accommodate feedback between doing and learning. This context could not.

Fitness for Execution was partial and fragile. Some teams implemented the mechanics and found them useful. The mechanics did not connect to anything above or outside the team. Execution fitness at the micro level is not nothing. It is also not sufficient.

Fitness for Improvement was a casualty. The feedback loops that existed were budget-cut casualties. The learning cycles that had begun to produce insight were terminated before they could inform any decision that mattered. The organization could not learn from what it was doing, because the systems that would have enabled that learning were treated as costs rather than capabilities.

The Impossible Premise

Here the question of what early diagnosis would have been worth becomes sharper. In the bank's case, there was at least a plausible path from correct diagnosis to different action, slim, but present. In the insurance company's case, the plausible path was narrower still.

The premise, agile without touching ways of working, was not a hypothesis that could be falsified by evidence. It was a prior commitment that determined what evidence would be permitted. An organization that has already decided what agile is and is not will receive a fitness diagnosis that contradicts that decision as a misunderstanding, not as a correction. The diagnosis would have been accurate. It would not have been usable.

This is the hardest thing the critical companion register requires saying: there are organizational situations where a correct approach cannot help, not because the approach is wrong, but because the conditions for using it are absent. The AFS is an approach that requires an organization capable of engaging with its own fitness. An organization that has decided in advance what its fitness means cannot do that. The instrument and the context are mismatched, and no amount of diagnostic precision closes that gap.

Micro Wins, Macro Illusions

In both cases, something real happened at the team level. This matters and deserves to be said plainly, without either overstating it or dismissing it.

In the bank, some teams developed genuine product ownership, genuine flow, genuine learning cycles. Individuals grew. Practitioners who had been running ceremonies without understanding them began to understand them. The vocabulary of fitness, fit for purpose, fit for context, fit for execution, gave people a language for describing what was working and what wasn't, which is a prerequisite for doing anything about it.

In the insurance company, the picture was similar and more constrained. Teams that participated by choice found the practices useful. They used agile disciplines to make visible the friction they had previously absorbed silently, the unclear priorities, the governance overhead, the dependency knots. The visibility did not produce organizational change. It produced individual clarity, which is not nothing, and is also not what the initiative was commissioned to achieve.

The AFS architecture would predict exactly this outcome. The diagnostic operates across four dimensions that are not independent. Fitness for Execution at the team level can improve even when Fitness for Purpose and Fitness for Context are broken above. Teams can become more adaptive within their own scope. What they cannot do, from within that scope, is repair the conditions that limit what their adaptiveness can accomplish.

Local fitness gains do not propagate upward into a system that is not designed to receive them. This is not a failure of the teams. It is a structural property of how organizations absorb — or fail to absorb — localized change. The micro wins were real. They were also contained. And containment, in both these cases, was the macro outcome.

What the AFS Would Have Said

Both cases fall clearly into what the Three Kinds of Wrong framework calls the third kind: a coherent approach used outside its valid domain. The approach itself was not wrong. The application was not wrong. The context was wrong, or more precisely, the context lacked the conditions that the approach requires to function at the system level.

The AFS approach/framework distinction is the relevant defence here. A framework, portable, repeatable, complete, would have travelled into both organizations unmodified and would have misapplied in both. It did, in the bank's case: the copied scaling framework was exactly this, and it produced exactly the result the distinction predicts. The AFS's insistence on practitioner judgment, on contextual reading before method selection, on fitness as the organizing question rather than maturity, these defences would have produced a more accurate assessment of what was possible in each context.

But accurate assessment is not sufficient. The AFS requires an organization capable of engaging with its own fitness. Both organizations, in different ways, were not. The bank could engage with it at the team level, and did. At the system level, the engagement was foreclosed by the decision to use a framework that had already answered the diagnostic questions the AFS would have asked. The insurance company could not engage with it at any level that mattered, because the premise of the engagement had pre-answered every question the fitness reading would have raised.

A correct AFS diagnosis, delivered early, would have named both of these conditions. It would have said: the purpose commitment here is rhetorical, not structural. It would have said: the context does not support system-level adaptation. It would have said: the expected outcome is micro-level gains that do not propagate. All of that would have been accurate. None of it would necessarily have been acted on.

What Good Diagnosis Cannot Do

The diagnostic instruments are good. The Four Questions produce a genuine reading of organizational fitness. The Adaptive Signal Field surfaces the dimensions that matter. The Friction-Fitness Map tests whether what is being called progress actually is. These are real capabilities, and they are worth the investment in developing them.

What they cannot do is substitute for the organizational conditions that make transformation possible. A fitness diagnosis is a reading of what is. It does not create what is not. An organization without genuine purpose commitment, without contextual readiness, without the willingness to build learning loops between doing and deciding, such an organization can be correctly diagnosed. It cannot be saved by its diagnosis.

This is uncomfortable to say, and worth saying anyway. The temptation in methodology writing is to imply that the right approach, correctly applied, produces the right outcome. It doesn't. The right approach, correctly applied to a context whose conditions are wrong, produces a correct reading of the conditions and the outcomes those conditions determine. That is a different and smaller thing.

The practitioners who worked these engagements did not fail because they lacked a better diagnostic instrument. They were working in conditions that constrained what any diagnostic instrument could accomplish. Better instruments would have named the conditions earlier and more precisely. They would not have changed the conditions.

There is a version of this paper that draws the conclusion that practitioners should screen engagements more carefully, should apply a pre-diagnostic before committing, should name the conditions explicitly before beginning, should reserve the right to decline when the conditions are absent. That conclusion is reasonable, and it is worth drawing. But it should not

obscure the harder one: that the conditions are often not fully legible at the outset. Purpose illusions look like purpose commitments until they don't. Impossible premises are often invisible until the intervention runs into them. The practitioner's judgment can improve the odds of reading the situation early. It cannot guarantee it.

Limitations

This retrospective is written from inside the engagements it describes. The author was the practitioner. That access produces texture that an external account cannot achieve. It also produces blind spots that an external account would not have. Both should be assumed.

The AFS diagnostic is applied here in retrospect, to situations that were not originally framed in AFS terms. The retrospective application is honest, the diagnostic logic was present in the practice even before the vocabulary was formalized, but it is not identical to a prospective application. A practitioner using the AFS at the outset of either engagement would have brought the structured vocabulary to the reading. That may have produced a more precise or a more communicable diagnosis. Whether it would have produced a different organizational response is a question this paper cannot answer.

The two cases described here are drawn from European financial services, banking and insurance. They share a regulatory culture, a governance orientation, and a risk posture that shapes what transformation looks like and what it is allowed to be. Practitioners in different sectors or different cultural contexts may find the failure patterns less familiar, or may recognize different versions of the same structural dynamics. The cases are specific. The argument is intended to be more general, with the caveat that more general arguments always lose some accuracy at the edges.

Finally, the micro-level gains described in both cases are real but undocumented beyond the practitioner's recollection. The organizational level outcomes, the macro failures, are clearer, because failures leave more visible traces than partial successes. The imbalance in evidence is worth naming.

Bridge

What changes after this paper?

Not the instruments. The Four Questions still do what they do. The Adaptive Signal Field still maps what it maps. The Friction-Fitness Map, the AFS Adaptive Engine, the four Level 2 capability domains, none of them are different on the other side of this retrospective.

What changes is the framing. The AFS is not a guarantee. It is a diagnostic capability, a way of reading organizational situations with more precision than the alternatives, producing more

honest assessments of what is possible and what is not. The precision is real. The honesty is real. The limit is also real: precision and honesty about conditions do not change the conditions.

The two cases here produced correct readings that could not be fully acted on. The micro-level gains were genuine. The macro-level outcomes were determined by conditions that pre-existed the diagnostic. That is a failure of context, not of approach. The distinction matters, not because it excuses anything, but because it locates the problem correctly. And locating the problem correctly is, in the end, what the AFS is for.

Even when it cannot fix what it has found.

Part of the AFS Open for Review series — adaptiveways.site/afs/open-for-review