


BEYOND THE CONTROL

The Regulated Sector Paradox: How Compliance Prevents the Change It Mandates

A Practitioner's Perspective on Compliance and Change in Regulated Industries



Giovanni Leonardi

July 2006

Contents

1. Executive Summary	3
2. The Shape of the Paradox	3
3. Why Compliance Crowds Out Change	4
4. The Machinery That Sustains It	5
5. What the Paradox Costs	6
6. Breaking the Pattern	6
7. The Longer View	7

The Regulated Sector Paradox: How Compliance Prevents the Change It Mandates

Transformation · No. 75

Originally written July 2006 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“Every control is a small vote against change, and an organisation that has accumulated ten thousand of them has, without ever deciding to, voted overwhelmingly to stay as it is.”

1. Executive Summary

Regulated organisations are told, in effect, to change. The rules that govern financial services, pharmaceuticals, utilities and the like are rarely content to freeze an industry in place; more often they demand that firms strengthen their controls, improve their transparency, protect their customers more carefully, and account for their conduct more fully than before. Change is the mandate. And yet the same organisations, held to that mandate, are frequently the least able to change of any in the economy. This is the paradox I want to examine: the machinery a firm builds to satisfy its obligation to change becomes, over time, the principal obstacle to changing at all.

The argument here is that the paradox is not an accident of poor management, nor a failure of will, but a structural consequence of how compliance is institutionalised. When an organisation converts a regulatory expectation into a permanent control, it freezes one particular answer to one particular question and defends that answer against all future revision. Do this ten thousand times and you have built an organisation optimised to preserve its own present state. The very discipline that lets a bank or an insurer demonstrate that it is safe is the discipline that makes it slow.

This essay traces the shape of the paradox, examines the structural forces that sustain it, counts what it costs, and closes with the beginnings of a resolution — not the abolition of control, which no serious practitioner would propose, but a different relationship between control and change in which the two are designed together rather than set against each other.

2. The Shape of the Paradox

Begin with the observation itself, because it is easy to state and easy to underestimate. Ask any executive in a heavily regulated firm whether their organisation finds change difficult, and they will agree at once. Ask them why, and the first answers will be the familiar ones: legacy systems, cultural inertia, the sheer size of the enterprise. Press a little further and a more interesting answer emerges. Change is difficult, they will eventually say, because everything has to go through compliance. Every new product, every altered process, every system migration must be checked against a body of rules, sign-offs and controls that has accumulated over years. The checking is not malicious. It is the organisation doing exactly what it was told to do. And it is precisely what makes the organisation slow.

Here is the paradox in its sharpest form. The regulator's intent is almost always progressive. Regulation of the modern kind – the wave that has intensified since the corporate scandals at the start of this decade, and that Sarbanes-Oxley made concrete for anyone listed in the United States – asks firms to become better: better governed, better controlled, more honest about risk. The regulator wants the firm to move toward something. But the instrument the regulator reaches for is the control, and a control is by its nature a device for preventing movement. It says: this shall not change without permission. Multiply that across a decade of rule-making and you have an organisation that has been asked, control by control, to hold still.

The regulator asks the firm to move toward a better state. The instrument the regulator reaches for – the control – is a device for preventing movement. The mandate is dynamic; the mechanism is static. The paradox lives in that gap.

I have watched this play out in the same way across very different institutions, which is what convinces me it is structural rather than particular. A firm receives a new obligation. It responds, sensibly, by designing a control: a checkpoint, an approval, a mandatory review, a piece of documentation that must exist before an action may proceed. The control works. The obligation is satisfied. The auditors are content. And then the control is never removed, because removing a control is frightening in a way that adding one is not. Over time the controls accrete like coral. Each is defensible on its own terms. Together they constitute a reef on which every attempt at change runs aground.

3. Why Compliance Crowds Out Change

Why does this happen so reliably? Several structural forces work together, and it is worth separating them.

The first is the asymmetry of blame. In a regulated firm, the person who adds an unnecessary control is almost never punished, while the person who removes a necessary one may end their career. The consequences of over-control are diffuse, slow and shared; the consequences of under-control are sharp, sudden and personal. Any rational actor facing that asymmetry will over-control. This is not cowardice. It is a correct reading of the incentives. And because it is correct, it is universal, which is why you find the same conservatism in institutions that share nothing else.

The second is the ratchet of documentation. Once a control is written down – in a policy, a procedure, a risk register – it acquires a kind of permanence that unwritten practice never has. To change it, someone must reopen the document, justify the change, secure approval, and accept ownership of the consequences. The written control is therefore stickier than the reasoning that produced it. Frequently the original reason has evaporated – the product it protected has been withdrawn, the risk it addressed has passed – but the control remains, because nobody is rewarded for the archaeology required to prove it obsolete.

- The control outlives the risk that justified it.
- The documentation outlives the control.
- The habit of deference to the documentation outlives everyone who understood why it was written.

The third force is the conflation of compliance with control, which I regard as the deepest of the three. Somewhere in the maturing of a regulated industry, a substitution takes place. The goal — being a well-run, honest, safe institution — is quietly replaced by a proxy: having the controls. The proxy is measurable, auditable and defensible, and the goal is none of these things, so the proxy wins. Once that substitution has happened, adding controls looks like virtue and questioning them looks like recklessness, regardless of whether any given control actually serves the goal. The organisation has stopped asking whether it is safe and started asking whether it can prove it followed the procedure.

“Every control is a small vote against change, and an organisation that has accumulated ten thousand of them has, without ever deciding to, voted overwhelmingly to stay as it is.”

4. The Machinery That Sustains It

If the forces above explain why controls accumulate, the machinery of the modern compliance function explains why they are so hard to dislodge. I want to be careful here, because it is fashionable to blame the compliance function itself, and that is both unfair and analytically lazy. The people who staff these functions are, in my experience, thoughtful and often acutely aware of the paradox they are caught in. The problem is the machine, not its operators.

Consider how the machine is built. A regulated firm typically organises its assurance in layers — the business that owns the risk, the compliance and risk functions that oversee it, and the internal audit function that checks the overseers. This architecture, which has become near-universal since the governance reforms of the past few years, is genuinely valuable. But it has an emergent property that nobody designed and everybody feels: each layer justifies its existence by finding things. An audit that finds nothing is suspected of not looking hard enough. A risk function that raises no concerns is thought to be asleep. The machine is therefore biased, at every level, toward the production of findings, and findings become controls, and controls accumulate.

What the machine optimises for	What the mandate actually needed
Demonstrable coverage of every risk	Intelligent attention to the risks that matter
Absence of audit findings	Presence of genuine safety
Adherence to the documented procedure	Achievement of the procedure's purpose
The ability to prove control after the fact	The ability to change safely in the moment

There is also the matter of tooling, and here the past few years have made things quietly worse even as they promised to make them better. The spread of enterprise systems and integrated risk platforms has allowed firms to encode their controls into software — into the workflow of an ERP implementation, into the mandatory fields of a case-management system, into the access rules of a data warehouse. Encoding a control in software is presented as modernisation, and in one sense it is. But a control embedded in a configured system is far harder to change than a control written in a manual, because changing it now

requires a technology project, a testing cycle and a release window. We have taken the stickiest thing in the organisation and set it in concrete. The firm that has automated its compliance has often automated its inability to change.

Finally, the machine sustains itself through language. The vocabulary of compliance is a vocabulary of permanence and prohibition — controls, requirements, mandatory, must, shall. It contains almost no words for revision, sunset, or graceful retirement. A profession that cannot easily say a thing in its own language will struggle to do that thing, and the compliance profession, for all its sophistication, has very little natural language for letting go.

5. What the Paradox Costs

It would be possible to treat all this as a grumble about bureaucracy, the sort of complaint every large organisation invites. That would be a mistake, because the costs are real and, more importantly, they fall in exactly the place the regulation was meant to protect.

The first cost is the slow death of good change. Every regulated firm has a graveyard of improvements that were never attempted because the compliance overhead made them uneconomic. A better way of serving a customer, a simpler process, a cleaner system — each is weighed not only against its own difficulty but against the tax of getting it through the machine, and many die at that second hurdle. The improvements that survive are the large, expensive, heavily-sponsored ones, because only they can absorb the overhead. The small, cheap, incremental improvements — the ones that in aggregate matter most — are precisely the ones the machine kills. The organisation loses its capacity for continuous small betterment and is left only with the capacity for occasional enormous upheaval, which is the worst possible distribution of change.

The second cost is the corrosion of judgement. When people work for years inside a system that rewards following the procedure over achieving the purpose, they gradually stop exercising the judgement the procedure was meant to encode. Why would they? Judgement is risky; the procedure is safe. Over a decade this hollows out the very capability the firm most needs in a crisis, when the procedures do not fit the situation and someone has to think. The regulated firm, having optimised for the demonstrable following of rules, discovers in its worst moment that it has trained a generation not to think for itself.

The third and subtlest cost is the illusion of safety. This is the one that should trouble the regulator most, because it is the exact opposite of what the regulation intended. An organisation buried in controls looks safe and feels safe, and the looking and the feeling are dangerous, because they discourage the harder question of whether it is safe. A firm can be in perfect compliance — every control operating, every document current, every audit passed — and be sliding toward failure, because compliance measures conformance to yesterday's understanding of risk, and risk does not stand still. The controls are a photograph of the dangers as they were understood when the controls were written. The world moves on; the photograph does not.

6. Breaking the Pattern

If the paradox is structural, then wishing it away is useless and moralising about it is worse. What is needed is a different structure. I do not have a tidy framework to offer — I

am suspicious of anyone who does, because a tidy framework would itself become another control — but I can describe the direction in which I have seen the paradox loosen, and the principles that seemed to matter.

The first principle is that controls must be born with an expiry date. The single most useful discipline I have encountered is the simple requirement that every new control carry, at the moment of its creation, an explicit statement of the risk it addresses and a date on which it will be reviewed for continued relevance. This does not remove the control. It merely denies it the automatic permanence that is the root of the accretion problem. A control that must periodically re-justify its existence is a control that can die a natural death, and a compliance estate in which controls die naturally does not turn into coral.

The second principle is that the firm must measure its control burden as deliberately as it measures its risk. Firms count their risks obsessively and their controls not at all — or rather, they count controls only as assets, never as liabilities. But a control is both. It mitigates a risk and it imposes a cost, and the cost includes the drag on change that this essay is about. An organisation that put the same rigour into understanding the total weight of its control estate as it puts into understanding its risk exposure would make very different decisions about which controls to keep.

1. Name the risk each control addresses, in writing, when the control is created.
2. Set a review date on which the control must re-earn its place or lapse.
3. Measure the aggregate burden of the control estate, not only its coverage.
4. Give someone explicit ownership of removing controls, with the standing and the incentives to do it.

That fourth point deserves emphasis, because it is the structural counter to the asymmetry of blame with which this essay began. The reason controls only ever accumulate is that adding is rewarded and removing is punished. The only durable fix is to create a role, with real authority, whose success is measured by the intelligent reduction of the control estate — someone for whom removing an obsolete control is a win rather than a risk. Until an organisation is willing to reward subtraction, it will drown in addition, and no amount of good intention will save it.

The deepest principle, though, is a change of mind rather than of mechanism. It is to stop treating compliance and change as opponents and start treating them as a single design problem. The firms that suffer the paradox most acutely are the ones that build their controls first and discover their rigidity later. The firms that suffer it least are the ones that ask, at the moment of designing any control, how will this control itself be changed when the world changes? — and build the answer in. A control designed to be revised is a fundamentally different object from a control designed to be permanent, even when the two do exactly the same thing today. The difference is entirely in whether the possibility of its own future change was part of its design.

7. The Longer View

I want to end by resisting the cynical reading of everything above, because it is available and it is wrong. The cynical reading is that regulation is the enemy of progress, that compliance is dead weight, and that the regulated firm is doomed to sclerosis. None of that follows. Regulation, at its best, is one of the few forces that reliably pushes firms to become better than their short-term incentives would make them. The mandate to change is a good mandate. The paradox is not that the mandate is wrong but that the standard instrument for delivering it — the accreting, permanent, defensively-held control — is

badly matched to it.

The practitioner's task, then, is not to choose between compliance and change, which would be a false and dangerous choice, but to dissolve the opposition between them. That means building compliance estates that can breathe: that add controls when risk demands and shed them when it recedes, that measure their own weight, that reward the people brave enough to simplify them, and that treat every control as a temporary answer to a permanent question rather than a permanent answer to a temporary one.

An organisation that learns to do this stops experiencing the paradox. Its compliance ceases to be the enemy of its transformation and becomes, instead, one of the disciplines through which it transforms safely. That is a harder thing to build than a wall of controls, and it will never look as reassuring in an audit. But it is the only version of a regulated firm that can actually do what its regulator, in the end, was always asking it to do: to keep changing, and to stay safe while it does.