


BEYOND SERIAL COMPLIANCE

Regulatory Intelligence Is a Capability, Not a Project

Why the serial, directive-by-directive response to overlapping regulation quietly costs a programme more than the standing capability it refuses to build — a view from financial services.



Giovanni Leonardi

March 2006

Regulatory Intelligence Is a Capability, Not a Project

Programme · No. 41

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

“The failure was that no one owned the space between the directives, which is exactly the space regulatory intelligence is meant to occupy.”

1. The same room, a different acronym

The programme director had seen the slide before. Not this exact slide — this one said MiFID readiness across the top, and eighteen months earlier it had said Basel II data foundations, and before that Sarbanes-Oxley 404 remediation — but the shape was identical. A RAG status turning amber. A dependency on client and instrument data at a level of granularity the firm did not hold in one place. A compliance lead with a thick binder of requirements, and an IT lead explaining, patiently, that the data existed but not in that form, not to that definition, and not by that date.

What made the meeting worth remembering was a single remark from the head of data, made almost to himself. We built most of this for Basel. We just built it to a different specification, and then we let the team go.

That sentence is the whole subject of this piece. Over the previous two years, financial firms had absorbed Sarbanes-Oxley, the first full year of reporting under IFRS, the long run-up to Basel II, the Third Money Laundering Directive, and now the approach of MiFID — each treated, funded and staffed as though it were the first regulation the firm had ever faced. It was not. The cost of pretending otherwise is a quiet, recurring tax, and this is a piece about refusing to keep paying it.

2. The mistake we make with every directive

Ask how an organisation handles regulatory change and you will usually be shown a project. It has a sponsor, a gap analysis, a remediation plan, a testing phase and a go-live. It is scoped tightly to the letter of the requirement, because it is funded as a cost and costs are minimised. It is staffed by pulling compliance and change people together for the duration. And when the deadline passes and the auditors sign off, it is disbanded, its people returned to their departments and its working knowledge dispersed with them.

Every part of that is locally rational and collectively ruinous. The tight scope forbids building anything durable, because durability looks like gold-plating to a cost sponsor. The temporary team means that the hard-won understanding of how this firm's data, systems and controls actually behave under regulatory pressure is never anyone's permanent possession. So the next directive arrives to an organisation that has, in the ways that matter, forgotten everything it learned from the last one.

The textbooks do describe a better shape. They draw a tidy cycle — scan, assess, implement, embed — and call it regulatory change management. What the diagram omits is that in most firms the scanning is done by no one in particular, the assessment happens

far too late to be cheap, and the embedding step is a box that is ticked but never actually performed, because the team that would embed anything has already been stood down.

3. Intelligence is not information, and it is not prediction

The word intelligence is doing real work here, and it is worth being precise about it. A subscription to a regulatory newsletter is information. A standing line on the risk register that reads forthcoming regulation is information. Neither is intelligence, because neither forms a judgement or changes a decision.

Regulatory intelligence, treated as a capability the programme owns rather than a document it files, has four attributes. Each is a muscle, and the firms that suffer least are simply the ones that have kept the muscle in use.

1. A horizon function that reads direction, not just text. It follows not only the published rule but the regulator's evolving posture — the FSA's steady drift toward principles-based supervision, and its Treating Customers Fairly work, were legible to anyone paying attention well before they hardened into specific obligations. Reading direction buys the one thing reactive compliance never has: time.
2. A translation function that turns a directive into implications for this firm and this plan. The generic requirement is worthless until someone can say what it means for our instrument data, our control environment, our delivery schedule. This is where most scanning efforts fail — they produce coverage, not consequence.
3. A memory, so that what was learned absorbing one wave survives to the next. The specifications, the data models, the control documentation, the hard-won map of where this firm's own information actually lives — these are assets, and disbanding the team that holds them is the equivalent of burning the drawings once the building is up.
4. A feedback loop into programme decisions early enough to be cheap. The cost of a regulatory change is not set at implementation. It is set eighteen months earlier, when an architecture was fixed without anyone in the room asking where regulation was heading.

That last point is the mechanism beneath all the others, and it deserves to be stated plainly.

The cost of complying with a regulation is largely decided before the regulation is even final — at the moment a programme fixes an architecture, a data model or an operating design with no one present who can see the regulatory direction of travel. By the time the requirement is published, the expensive decisions have already been made.

4. What it actually costs — a composite

Consider a firm — a composite, but nothing in it is unusual — that ran two programmes in parallel across 2005. One delivered the transition to IFRS reporting; the other built the data foundations for a Basel II internal-ratings submission. They had different sponsors, different steering committees and different advisers, and they never spoke, because on the surface one was an accounting change and the other a credit-risk change.

Both, underneath, depended on the same thing: transaction and exposure data reconciled to the general ledger at a finer granularity than the firm had ever routinely maintained. Neither programme specified it the same way. The IFRS team built its view first, to its own definitions. The Basel team, arriving months later, found that view unusable for its purpose and rebuilt the underlying extract from scratch.

The direct rework ran to a few hundred thousand pounds of effort — irritating, but survivable. The real cost was nine months. The rebuilt data pushed the internal-ratings submission back by three quarters, which delayed the capital benefit on which the whole Basel programme had been justified. A single regulatory-intelligence view — one function whose job was to hold both directives in mind at once — would have seen the shared dependency at the point the IFRS data model was being drawn, when accommodating it would have cost a design conversation rather than a second build.

Nobody in this story was incompetent. Each programme did its own job well. The failure was that no one owned the space between the directives, which is exactly the space regulatory intelligence is meant to occupy.

5. The honest objection: you cannot forecast a regulator

The strongest argument against everything above is not foolish, and it deserves its best form rather than a caricature. It runs like this. Regulation is genuinely unpredictable. Directives are amended in committee, timetables slip — MiFID's own deadline had only just moved — and priorities shift with the political weather. You cannot build reliable intelligence about something you cannot foresee. A permanent regulatory-intelligence function is therefore an expensive standing army that will spend most of its time guessing wrong, and the rational strategy is to pay each bill as it arrives: staff up sharply when a requirement becomes concrete, stand down when it is met.

The objection lands cleanly on one point and misses on another. It is right that you cannot forecast the specific rule. It is wrong to conclude that intelligence is therefore prediction. Intelligence is not the ability to name next year's directive; it is the possession of the muscles that every directive turns out to need — reconciled data at defensible granularity, a documented and genuinely understood control environment, senior attention that can be mobilised without a standing start. The firms that came through Sarbanes-Oxley most cheaply were not the ones who had predicted Section 404. They were the ones who already understood their own controls, and so had less to discover under time pressure. What the capability produces is not a forecast. It is optionality — the ability to absorb whatever arrives at a lower marginal cost, because the foundations are already there and already remembered.

And the comparison the objection relies on is the wrong one. It weighs the standing function against nothing, when the real alternative is the serial rebuild — the same data foundation specified and discarded three times, the same understanding reassembled from a cold start by a new team each cycle. Measured against that, the standing capability is not the expensive option.

“Intelligence is not the ability to name next year's directive; it is the possession of the muscles that every directive turns out to need.”

6. The uncomfortable part of the account

If the case is this clear, why does almost every organisation still choose the serial project? Because an honest account has to admit that the project is locally rational for the people making the choice. Regulation is funded as cost, and a permanent capability is a permanent line item that someone must defend every budget round, whereas a project ends and stops costing money. The reward structure pays for passing the audit, not for being cheaper to comply next time — no one is thanked today for a data decision that will save nine months in two years' time. And the credit for a disaster quietly averted is invisible by construction. These are not stupidities to be scolded away; they are precisely why the pattern is so stable.

There is a genuine risk on the other side, too, and pretending otherwise would be its own dishonesty. A standing regulatory-intelligence function can ossify into exactly the bureaucracy its critics fear — a horizon-scanning unit that generates thick monthly reports no decision-maker reads, coverage mistaken for consequence. The discipline that prevents this is easy to state and hard to hold: the capability is justified only by the decisions it changes, never by the completeness of what it monitors. If it is not altering the design of programmes early enough to be cheap, it has become the very cost the sceptic warned about.

So the account ends without triumph. Most programmes will go on treating each directive as a shock, because the funding and reward structures make that the path of least resistance. The practitioners who break the pattern are rarely the ones who win the argument in a steering committee. They are the ones who, handed the current directive, quietly use it as the occasion to build the muscle rather than merely to pass the test — who specify the data foundation once, properly, and refuse to let the team that understands it simply disperse. The next wave is already forming somewhere in a consultation paper. The only real question is whether we will meet it having remembered anything at all.