

BEYOND BLANKET CAUTION

AI in Regulated Industries — Moving First When You're Built to Move Last

Turning Model Risk Heritage into the Regulated Sector's AI Deployment Advantage

Giovanni Leonardi

March 2025



AI in Regulated Industries — Moving First When You're Built to Move Last

Transformation · No. 410

March 2025

"Waiting for a rulebook that the regulator has no intention of writing is not caution; it is misreading the regulatory compact."

1. The Asymmetric Bet Nobody Wants to Take

The conversation in most regulated boardrooms follows a script by now. Someone presents the AI opportunity — the productivity gains, the competitive threat from less constrained entrants, the consultancy slides showing exponential curves. Then the Chief Risk Officer speaks. The discussion turns to model validation, regulatory expectations, conduct risk, and the meeting concludes with agreement to "proceed carefully," which in practice means proceeding barely at all.

The script is understandable. Financial services firms, insurers, and utilities are built to punish error. Their cultures, governance structures, and incentive systems have been shaped by decades of regulatory consequence — from enforcement actions to licence revocations. The instinct to wait until the regulatory picture clarifies is not irrational. It is, however, wrong.

It is wrong because it mistakes the nature of the advantage these firms hold. The prevailing assumption — that regulation is a brake on AI adoption — inverts the actual dynamics. Regulated industries possess something their less constrained competitors lack and cannot easily acquire: a mature infrastructure for governing models under uncertainty. The firms that recognise this and deploy it will not merely keep pace with unregulated entrants; they will surpass them, because they will have solved the hardest problem in enterprise AI — trust at scale — before anyone else arrives at the question.

2. The Topology That Matters

The resolution to the regulated firm's AI dilemma is not a choice between caution and courage. It is topology.

Most organisations that describe themselves as "proceeding carefully" with AI have not done the foundational work of mapping their own processes according to where regulatory teeth actually bite. Instead, they have applied a blanket caution born of an undifferentiated sense that everything they do is sensitive. This is the category error that

paralyses.

In my experience across financial services transformation, the proportion of processes that carry genuine regulatory consequence — where an AI failure could trigger enforcement, customer harm, or prudential risk — is substantially smaller than organisations assume. In a typical retail bank, perhaps fifteen to twenty per cent of operational processes sit in this category: credit decisioning, financial promotions, complaints handling that touches conduct obligations, and the thin layer of genuinely regulated advice. These are the processes where the PRA's expectations under SS1/23, or the EU AI Act's high-risk classification, impose real constraints on how models can be developed, validated, and deployed.

The remaining eighty per cent — internal research and analysis, drafting and communications preparation, code generation, operational workflow, service preparation, management reporting — are processes that feel sensitive because they happen inside a regulated firm, but carry no more regulatory consequence than the same activity performed inside a technology company. The discomfort is inherited from the institution's identity, not derived from the regulatory framework.

This distinction — between processes with teeth and processes that merely feel sensitive — is the topology that matters. And mapping it honestly is the single most productive thing a regulated firm can do to unlock AI deployment at scale.

3. Deploying Hard into the Eighty Per Cent

Once the topology is drawn, the strategic implication is immediate. The genuinely low-stakes majority of processes — the eighty per cent — should be the site of aggressive, rapid AI deployment. Not pilots. Not proofs of concept. Full operational deployment, with the same urgency a technology firm would bring.

We have been slower to reach this conclusion than we should have been, in part because the governance apparatus of regulated firms does not distinguish well between "this requires regulatory-grade validation" and "this happens inside our building." The model risk management framework inherited from Basel — now codified in the PRA's SS1/23 — was designed for a world where "model" meant a credit scoring algorithm or a pricing engine: something with direct financial consequence, deployed into production with clear inputs and outputs. Applying the same validation overhead to an internal summarisation tool or a code-generation assistant is not rigour; it is misallocation.

The firms that have moved fastest have done something deceptively simple. They have created an explicit tiering within their model risk frameworks — not abandoning governance, but calibrating it. A large-language-model application that drafts internal briefing notes receives lighter-touch oversight than one that generates customer-facing communications, which in turn receives lighter oversight than one that influences a regulated decision. The tiers are defined not by the technology but by the regulatory and conduct consequence of the output.

This is not a radical proposition. It is what proportionate risk management has always meant. But the psychological barrier in regulated firms is real: the fear that any AI deployment that goes wrong will be viewed through the lens of the firm's regulated

status, regardless of whether the specific use case carried regulatory weight. Overcoming that fear requires the topology to be drawn explicitly, endorsed at board level, and communicated to the regulator — not as a request for permission, but as a demonstration of the firm's own risk framework in action.

4. The Heritage Advantage

Here is the argument that the caution narrative entirely misses: the model risk management infrastructure that regulated firms have spent fifteen years building is not an obstacle to AI adoption. It is a competitive advantage of extraordinary value.

Consider what an unregulated technology firm must build from scratch when it discovers — as it inevitably does — that its AI deployments carry risks that require governance. It must establish validation frameworks, create model inventories, define escalation paths, build monitoring capabilities, and develop the organisational muscle to say "no" to a deployment that fails its own standards. This is precisely the infrastructure that banks, insurers, and asset managers already possess.

The PRA's SS1/23 framework, for all the compliance burden it imposes, has given UK financial services firms something genuinely valuable: a tested methodology for distinguishing between models that work and models that merely appear to. The disciplines of model validation — independent challenge, out-of-sample testing, ongoing performance monitoring, clear accountability for model owners — are exactly what responsible AI deployment demands. The regulated sector does not need to invent these disciplines; it needs to extend them.

The constraint, taken early, becomes the moat. Regulated firms that treat compliance as a design partner — building validation, monitoring, and accountability into their AI deployments from the first use case — will find themselves with a governance infrastructure that unregulated competitors spend years trying to retrofit.

The EU AI Act reinforces this dynamic. Its risk-based classification system — distinguishing between unacceptable, high, and limited-risk AI applications — mirrors the topology approach. Firms that have already mapped their processes against regulatory consequence will find the Act's requirements not an additional burden but a confirmation of work already done. The conformity assessments, the risk management systems, the human oversight requirements: these are extensions of existing model risk management, not novel impositions.

5. The Strongest Objection — and Why It Falls Short

The serious counterargument is not that regulated firms should avoid AI. It is that the regulatory environment is still moving, and early deployment risks building on foundations that regulators later reject. Better, this argument runs, to wait for definitive guidance and build once, correctly.

This objection has surface appeal but fails on two counts. First, regulators — the PRA, the FCA, the European Supervisory Authorities — have been notably clear that they expect firms to develop their own frameworks rather than wait for prescriptive rules. The UK's approach, deliberately sector-specific rather than legislative, is an explicit signal that the regulator wants to see mature, proportionate risk management from the firms themselves. Waiting for a rulebook that the regulator has no intention of writing is not caution; it is misreading the regulatory compact.

Second, and more fundamentally, the firms that deploy early into the low-stakes eighty per cent are not taking regulatory risk. They are building operational experience, training their validation teams on new model types, developing monitoring approaches, and accumulating the evidence base that regulators will demand when the time comes to deploy into the consequential twenty per cent. The early mover does not build on foundations that might shift; the early mover is building the foundations, in the space where the regulatory stakes allow for learning.

6. The Real Risk Is Standing Still

The asymmetric bet facing regulated industries is not the one most boardrooms assume. The risk is not that early AI adoption triggers regulatory consequence. The risk is that delayed adoption cedes operational efficiency, talent, and ultimately market position to less constrained entrants — while simultaneously failing to develop the institutional capability that genuinely consequential AI deployment will require.

We have been here before. The regulated sector's response to cloud computing followed the same pattern — years of hesitation driven by data-sovereignty concerns and an undifferentiated sense of regulatory risk, followed by rapid adoption once it became clear that the competitive cost of delay exceeded the compliance risk of action. The firms that moved first into cloud did not do so recklessly; they mapped the data, tiered the workloads, and engaged their regulators. They moved with the topology, not against it.

The same discipline applies now. Map the processes. Draw the line between teeth and feel. Deploy hard into the space where the regulatory stakes are genuinely low, and use that deployment to build the governance muscle for what comes next. The regulated firm's heritage is not a handicap. It is, for those willing to recognise it, exactly the advantage the moment demands.