

BEYOND THE MODEL GAP

The AI-Native Competitor — What Incumbents Should Actually Fear (and What They Shouldn't)

The Coordination-Cost Gap That No AI Programme Can Close

Giovanni Leonardi

December 2025



The AI-Native Competitor — What Incumbents Should Actually Fear (and What They Shouldn't)

Transformation · No. 426
December 2025

"The assets of incumbency are a window, and a window is only an asset if what happens inside it is restructuring."

1. The Boardroom Conversation We Keep Having Wrong

The scenario surfaces in strategy sessions with remarkable consistency now. A board member or a nervous executive raises the spectre: a startup, founded this year, building something close to our core proposition with a fraction of the headcount. The conversation that follows is almost always about the model — which large language model they are using, whether their AI is better than ours, how much more we should be investing in our own AI capabilities. The implied strategic question is about a technology gap.

I believe we are looking at the wrong gap.

The AI-native firms beginning to appear across financial services, professional services, logistics, and insurance do not, in the main, possess superior artificial intelligence. The models they use are the same models available to everyone — rented by the month, upgraded by the provider, interchangeable with the next generation. What these firms possess is the absence of something: the accumulated coordination cost that large organisations carry as a structural tax on everything they do.

No legacy technology estate requiring integration programmes before a single workflow can change. No functional silos demanding months of cross-departmental alignment before a process can be redesigned. No enterprise architecture review boards governing what connects to what. No change advisory boards sequencing deployments across shared infrastructure. Their advantage is not that their AI is better. It is that their AI has almost nothing to be retrofitted around.

This distinction matters enormously for how incumbents respond. And we are not yet, as a profession, thinking about it with anything like the precision it demands.

2. The Model Gap Is Not the Threat

Let me be direct about the part of this that is less threatening than it appears.

The large language models that underpin most current AI applications are, as of late 2025, increasingly commoditised. They are rentable. A firm can access frontier-class capabilities — reasoning, code generation, document analysis, conversational interaction — through API subscriptions costing a fraction of what custom model development required even two years ago. The performance gap between what a well-resourced incumbent can deploy and what a lean startup can deploy, in terms of raw model capability, is narrower than at any point in the history of enterprise technology.

This is, counterintuitively, reassuring for incumbents — or it should be. The nightmare scenario in which a competitor possesses fundamentally better intelligence is not materialising. What is materialising is something more subtle and, I suspect, more consequential.

3. The Operating-Model Gap

The coordination costs embedded in a large organisation are not incidental. They are structural. They accumulated for good reasons — regulatory compliance, risk management, operational resilience, the governance demands of scale. But they are real, they are heavy, and AI is making them newly visible because AI, for the first time, offers a credible alternative to many of the activities they exist to coordinate.

Consider a composite but representative example. A mid-tier insurance firm decides to automate a claims-handling workflow. The model selection takes a fortnight. The integration work — connecting the model to the policy administration system, the document management platform, the compliance engine, the downstream payment system — takes nine months. The operating procedure redesign, the retraining, the union consultation, the change advisory approvals, the parallel-run period: another six months. Total elapsed time from decision to live capability: approximately fourteen months for a single workflow. Total cost attributable not to the AI itself but to the organisational overhead of deploying it: roughly eighty per cent of the programme spend.

An AI-native competitor building the same capability from scratch faces none of this. The workflow is designed around the model from the start. There is no legacy system to integrate with, no parallel-run period because there is no prior process, no organisational change programme because the organisation was built this way. What takes the incumbent fourteen months takes the native firm, perhaps, eight weeks.

The gap is not in the intelligence. It is in the cost of change. And this is the gap that no AI programme, however well funded, can close — because the coordination costs are not bugs in the incumbent's operating model. They are the operating model.

The strategic question is not which model to invest in but which coordination costs to dismantle — because those costs are the entire competitive surface an AI-native entrant needs.

4. A Cost-of-Coordination Audit

If the diagnosis is structural rather than technological, the strategic response must be structural too. The move that matters is not another AI programme — it is what I would describe as a cost-of-coordination audit: a disciplined examination of where the organisation's own structure taxes it most heavily, because that tax is the entire opening an AI-native competitor needs.

The audit asks a specific set of questions, and they are not the questions most AI strategies are currently asking:

1. Where does the organisation spend the most elapsed time and cost not on work, but on the coordination of work — handoffs, approvals, alignment meetings, integration testing, change management?
2. Which of those coordination activities exist because of structural choices (functional silos, shared platforms, centralised governance) rather than because of genuine regulatory or risk requirements?
3. For each structural coordination cost identified, what would a competitor look like that simply did not carry it — and how would that competitor's unit economics differ?
4. Which coordination costs could be restructured within the next eighteen months, and which are genuinely load-bearing — so deeply embedded in the regulatory or operational fabric that removing them would create unacceptable risk?

The fourth question is as important as the first three. Not all coordination cost is waste. Some of it is the price of operating at scale in a regulated environment, and a serious audit distinguishes the load-bearing from the parasitic. The AI-native competitor's advantage is not infinite — it exists precisely in the space where coordination cost is structural but not necessary.

The output of this audit is not a technology roadmap. It is an operating-model roadmap: a sequenced plan to restructure the organisation's costliest coordination points, starting with those most exposed to native competition and most amenable to change.

5. What Incumbency Still Buys

It would be a mistake to read this analysis as a counsel of despair. Incumbents possess assets that AI-native competitors cannot replicate quickly, if at all.

Trust and regulatory standing. In financial services, insurance, healthcare, and critical infrastructure, the licence to operate is not merely legal — it is reputational. Customers, regulators, and counterparties extend a degree of trust to established institutions that a new entrant must earn over years. This is a genuine and durable advantage.

Distribution. Established client relationships, channel partnerships, and embedded positions in industry workflows give incumbents access to demand that a new entrant must build from zero.

Data. Decades of transactional, operational, and behavioural data — often messy, often siloed, but real and extensive — represent a substrate that, once properly accessible,

makes AI applications substantially more effective than the same models trained on public data alone.

Institutional knowledge. The tacit understanding of how a market actually works, where the exceptions cluster, what the edge cases look like — this is not captured in any model and not available at any price to a new entrant.

These assets are real. But they buy time, not immunity. The coordination-cost gap is not static; it compounds. Every quarter an incumbent spends fourteen months deploying what a native firm deploys in eight weeks, the gap in delivered capability widens. The assets of incumbency are a window, and a window is only an asset if what happens inside it is restructuring.

6. Where This Could Break

A foresight piece that does not examine its own assumptions is not foresight but advocacy. Several forces could alter the trajectory I have described.

It is possible that the coordination costs I have characterised as structural prove more malleable than expected. If the next generation of enterprise integration platforms dramatically reduces the cost of connecting legacy systems — and there are credible efforts underway — then the incumbent's deployment overhead shrinks and the native competitor's structural advantage narrows. We are seeing early signs of this in middleware tooling, though whether it scales to genuine enterprise complexity remains, I think, an open question.

It is also possible that AI-native firms encounter their own coordination costs as they grow. Scale has a way of imposing governance, and regulatory environments have a way of imposing process. The lean, agent-first operating model that works at fifty employees may prove difficult to sustain at five hundred, and some of the coordination overhead incumbents carry may turn out to be not the legacy of poor decisions but the irreducible cost of operating at scale. If so, the competitive advantage is transient — real but self-limiting.

And it is possible that the competitive dynamics play out not as displacement but as acquisition — incumbents purchasing native firms for their operating models rather than their technology, grafting the coordination-light structure onto the incumbent's assets. Whether this works in practice, or whether the organisational antibodies of a large firm simply reject the transplant, is something we will only learn through experience.

What I do not think is likely is that the coordination-cost gap proves irrelevant. The forces driving it — the commoditisation of AI capability, the structural weight of legacy operating models, the sub-linear scaling of AI-native workflows — are too fundamental to be resolved by incremental improvement. The question is not whether the gap matters but how incumbents choose to spend the time their existing assets buy them.

The cost-of-coordination audit is not the whole answer. But it is, I believe, the right place to start — because it asks the question the AI programme alone never does: not what can we automate? but what in our own structure is the opening our competitors need?