

Beyond the Prompt

The AI-Native Firm

LinkedIn and kjerside.com article | AI series, Part 3 of 5 | 29 July 2026

Series note: This is the third article in a five-part series on AI, economics and advisory work.

The prompt was a useful gateway into AI. It taught professionals that a model could draft, summarise, code, translate and increasingly reason through multi-step problems. But the prompt is too small a unit of analysis.

A prompt is a request. A firm is a decision system. Between the two sits the real frontier: the design of cognitive production.

The prompt was the demo. The workflow is the strategy.

The first article in this sequence argued that AI changes the scarcity economics of cognition. The second argued that intelligence is layered and often collective.

This article takes the next step. If

intelligence can be reproduced in parts, the strategic question is not only which model to use. It is how to organise models, people, memory, challenge and decision rights into a production system.

That is what I mean by an AI-native firm. Not a normal firm with chatbots attached. A firm whose core analytical workflows have been redesigned around the new economics of cognition.

The companion essays unpack the technical layers beneath that operating model. This article stays at the organisational level: how should those layers be arranged into recurring work?

A prompt is a request. A firm is a decision system. Between the two sits the real frontier: the design of cognitive production.

1. A prompt is useful, but it is not an operating model

Most organisations still approach AI as a tool adoption problem: give people access, run training, encourage prompt libraries, measure usage. This is a rational starting point, but it keeps the old organisation intact. The model becomes an assistant inside workflows designed for a different cost structure. Most firms are automating the employee while preserving the bureaucracy.

That leaves much of the value on the table. In professional work, the expensive part is rarely one answer in isolation. It is the sequence: defining the question, gathering evidence, building a view, challenging assumptions, reconciling disagreement, preparing a recommendation and carrying it into a decision.

Most firms are automating the employee while preserving the bureaucracy.

In M&A, for example, diligence is not one clever question. It is a structured chain of judgement: investment thesis, market structure, customer evidence, competitive dynamics, financial model, management claims, risk register, value-creation plan and board recommendation. A model can help at many points in that chain. But the value appears only when the chain itself is redesigned.

Reasoning models make this more important, not less. If a model can plan, debug and synthesise more effectively, the scarce question is no longer whether it can produce an answer. It is whether the answer sits inside the right workflow: with the right evidence, assumptions, constraints, review gates and decision owner.

Better reasoning raises the return to better operating design.

2. The real unit is the cognitive workflow

A cognitive workflow is the path by which ambiguity becomes action. It is not a chain of prompts. It is the set of roles, hand-offs, memory, review gates, tools, verification standards and decision rights that turn raw information into a recommendation someone is willing to own.

This matters because AI changes the cost and speed of several steps at once. Research can be compressed. First drafts arrive faster. Alternative hypotheses can be generated in parallel. Weak arguments can be challenged before they reach a meeting.

The bottleneck moves from production capacity to workflow quality.

The organisation that merely accelerates old workflows may simply produce more noise. The organisation that redesigns the workflow can change the economics of the work itself: smaller teams, faster cycles, sharper challenge, more explicit assumptions and clearer escalation. That is the difference between AI-enabled and AI-native. AI-enabled firms add models to existing workflows. AI-native firms redesign the workflow around the fact that parts of cognition have become cheaper, faster and more scalable.

3. AI-native workflows need roles, memory and challenge

Once intelligence is understood as partly collective, the frontier shifts from one model to a system of differentiated roles. The core question is how to build productive tension without recreating the cost and friction of the old organisation.

A serious AI-native workflow should not ask one generic assistant to do everything. It should separate functions: research, analysis, modelling, scepticism, synthesis, editing, risk review and decision support. These do not have to be literal agents in every case. The important point is functional separation. Different roles create productive tension.

The system also needs active memory and shared language. Without memory, it repeats work and loses context. Without shared definitions, it becomes verbally fluent but conceptually loose. A long context window is working material, not memory. Memory in the institutional sense is continuity under governance: what is retained, for whom, under what permission, and with what audit trail.

Finally, the workflow needs challenge. The first plausible answer should not be accepted by default. Traditional firms reduced error through hierarchy, peer review and team tension. AI-native workflows need their own equivalents.

There is more design beneath this: a set of working principles for role separation, memory, verification, challenge, synthesis and escalation on which I build AI-native workflows and companies. The principles travel. The detail is the craft, and it stays in the workshop.

4. Error management is the real test

The common objection is that AI makes mistakes. It does. But that objection is incomplete because humans also make mistakes. The question is not whether a production system has errors. Every production system does. The question is how errors are surfaced, challenged and contained.

This is why the operating model matters as much as the model. A weaker model inside a stronger workflow may outperform a stronger model used casually. Conversely, a powerful reasoning model inside a loose workflow can create confident noise at scale - or optimise efficiently towards the wrong objective. A model can be brilliant and the firm can still be incapable.

In an AI-native firm, error management should be designed rather than hoped for. Assumptions should be visible. High-stakes outputs should be challenged. Decisions should have clear owners. The system should know when not to decide.

A model can be brilliant and the firm can still be incapable.

This is not bureaucracy for its own sake. It is the digital equivalent of professional discipline. Consulting firms, law firms, investment committees and audit processes all evolved review mechanisms because high-value cognition carries consequences. AI-native firms need them for the same reason.

The new failure mode is not only hallucination. It is mis-specified agency: the system pursues a plausible but wrong objective, treats a soft constraint as binding or a hard constraint as optional, or acts on stale and partial context. The response is not better prompting. It is objective-setting, constraint management, verification and escalation by design.

5. From AI-enabled incumbents to AI-native firms

In the short run, most firms will treat AI as an assistant inside old structures. That is understandable. Incumbents protect the existing organisation first and redesign it second. But the deeper gains will not come from scattered tool use. They will come from rethinking workflows, staffing models, management layers and the hand-offs between problem definition, analysis, review and decision.

The sectors most exposed are those where valuable work is already codified in language, analysis, software and document production: law, finance, research, software, media and large parts of consulting. That does not mean these professions disappear. It means their internal economics change.

AI does not replace professions as wholes. Professions are bundles of tasks, relationships, responsibilities and accountabilities. AI reprices tasks inside those bundles.

Professional arena	Work becoming cheaper	Scarce residual work
Consulting and strategy	Market scans, issue trees, synthesis, first-draft analysis, benchmark gathering.	Problem selection, client trust, political judgement, decision facilitation, accountability.
PE and M&A	Data-room review, diligence Q&A, risk registers, market mapping, financial-model checks.	Investment judgement, management assessment, negotiation, risk appetite, IC conviction.
Law	Contract review, precedent search, clause comparison, drafting.	Legal strategy, negotiation, fiduciary duty, client accountability.
Software	Code generation, debugging, test scaffolding, refactoring.	Architecture, security responsibility, product judgement, live-system ownership.
Corporate functions	Policy Q&A, process documentation, reporting, workflow automation.	Operating-model design, escalation, incentives, cross-functional alignment.

The exposed unit is not the job title. It is the workflow. Fewer people may be needed for some stages, more output expected from others, and the premium shifts toward problem selection, client trust, orchestration and accountable judgement.

Over time, this suggests the emergence of AI-native firms rather than merely AI-enabled incumbents. Once the production model changes enough, the organisational form changes with it. For advisory work, that possibility is especially uncomfortable and especially interesting. Consulting firms were built around elite cognition, structured problem solving and team-based challenge.

Those disciplines still matter. But they may no longer require the same organisation, staffing model or cost structure.

Closing thought

The future of AI in professional work will not be won by the cleverest prompt. Prompts matter, but they are not strategy. The strategic unit is the designed path from ambiguity to action.

As AI makes more preparatory cognition abundant, the scarce work moves to framing, verification, orchestration, trust and accountable judgement. The firms that benefit most will not simply attach models to old workflows. They will redesign the workflow itself.

The next article moves from operating model to application: what it means to rebuild parts of consulting from first principles around AI-native workflows, while preserving the disciplines that made advisory work valuable in the first place - structure, challenge, synthesis, judgement and accountability.

Disclosure

This article was AI-assisted. The argument, judgement and conclusions are my own.

Jesper Kjerside is co-founder of *taonta*, an AI-native advisory firm, and *fastland*, the software company behind it.

AI series, Part 3 of 5 | 29 July 2026