

The New Boundary of the Firm

Why AI Will Not Dissolve Companies, But Redefine What They Are For

A regime-conditioned generalization of Coase for the AI era

Salim Ismail & Ted Shelton

Executive Summary

For nearly a century, business leaders have used a simple question to decide what belongs inside the company and what should be bought from the market: is it cheaper and safer to coordinate this activity ourselves, or to contract with someone else to do it?

Digital platforms, cloud infrastructure, APIs, and now AI agents have changed the underlying economics of coordination. Markets have become easier to search, monitor, and manage. External partners can be found, connected, measured, and replaced faster than ever. At the same time, AI can now perform work that once required layers of human coordination inside the firm. The result is not the end of the company. It is the beginning of a different kind of company.

The firm is moving from being a container for human coordination to being a container for trusted agency.

We are not disputing the core hypothesis that Coase proposed, that firm boundaries are determined by comparative coordination costs; what has changed is the technology that defines those costs. This paper therefore does not reject Coase. It updates the Coasean boundary condition for a world of digital platforms, AI agents, and institutional accountability.

This shift has five major implications for business leaders.

First, digital infrastructure has made company boundaries more flexible. The old make-or-buy decision was often anchored in transaction cost. Today, the decision often turns on speed, control, data access, flexibility, and the ability to change direction.

Second, AI does not automatically favor outsourcing. It can reduce coordination costs inside the company and outside it. What matters is where the company can place trust. If a company can trust its internal AI more than external AI services, it may expand its scope while reducing headcount. If external AI markets are more capable, trusted, and standardized, the company may outsource more.

Third, the most important strategic question is not simply "Can AI do the task?" It is "Can we safely let AI act without a human checking every step?" The economic value of AI changes dramatically when it moves from assisting humans to replacing parts of the human coordination chain.

Fourth, accountability becomes more important as technical coordination gets cheaper. AI agents cannot hold licenses, accept fiduciary duties, bear reputational consequences, or pay damages. Legal and commercial responsibility still has to attach to an institution. This is why firms persist even when software can coordinate work at enormous scale.

Fifth, in high-stakes sectors, the logic of outsourcing may reverse. When external AI failures can cascade through financial systems, healthcare networks, mobility platforms, defense infrastructure, or regulated

data environments, large firms may need to internalize more, not less. They will do so not because hierarchy is cheaper, but because accountability must be contained.

The strategic challenge for executives is therefore not to make the company as large or as lean as possible. It is to decide which activities require internal trusted agency, which can be safely delegated to external markets, and which must remain under institutional control because the liability of failure is too large to outsource.

1. The Question Every Company Has To Answer

Every business is built around a boundary.

Inside the boundary are the people, assets, systems, data, processes, and responsibilities the company chooses to own or directly control. Outside the boundary are suppliers, contractors, partners, platforms, marketplaces, and customers. Strategy is partly the art of deciding where that boundary should sit.

Should a retailer own its logistics network or rely on third-party carriers? Should a bank build its own AI risk models or buy them from a vendor? Should a healthcare company run its own data infrastructure or use a cloud platform? Should a manufacturer vertically integrate critical components or coordinate a network of specialist suppliers?

These are not merely operating decisions. They determine what the company is.

Ronald Coase offered the classic explanation for why firms exist at all. Markets are powerful, but they are not free to use. Finding suppliers, negotiating contracts, monitoring performance, enforcing agreements, and resolving disputes all consume time and attention. When the cost of using the market becomes higher than the cost of directing work internally, companies bring activities inside the firm.

That insight remains sound. But it was developed in a world where information moved slowly, search was expensive, monitoring was episodic, and most coordination depended on human judgment. The modern business environment is different. Digital platforms have made market coordination cheaper and faster. AI is now doing the same for internal coordination. The old question remains, but the machinery underneath it has changed.

The practical question for leaders is no longer simply whether hierarchy is cheaper than the market. It is whether the company has the data, trust, accountability, and governance needed to let intelligent systems act on its behalf.

2. Digital Platforms Made Boundaries More Elastic

The first major change came before generative AI. It came from the digitization of coordination.

Cloud computing, APIs, platform marketplaces, real-time telemetry, digital payments, identity systems, reputation scores, and standardized data interfaces all reduced the friction of working across company boundaries. Search became faster. Vendor onboarding became easier. Monitoring became continuous. A partner's performance could be measured in real time rather than inferred from quarterly reviews. Software made it possible to coordinate far-flung networks without owning every participant in the network.

This is why digital-native businesses often look less vertically integrated than industrial-era firms. They can plug into payment processors, logistics networks, cloud providers, software vendors, creator platforms, and data services instead of building every capability themselves.

But digitization did not simply push every company toward outsourcing. It made boundaries more sensitive.

When market coordination becomes cheap, other strategic factors move to the foreground. Speed can matter more than nominal cost. Control over data can matter more than vendor pricing. The ability to adapt can matter more than operating efficiency. A company may outsource when flexibility is paramount, then reintegrate when latency, quality, or customer experience becomes decisive.

Amazon's logistics strategy illustrates the point. The company did not build logistics capabilities merely because outside providers were unavailable. It built them because speed, control, data, and customer promise became central to the business model. In a digital economy, the question is often not "What is the lowest-cost provider?" but "Where must we control the feedback loop?"

Digital technology therefore made firm boundaries more elastic. Companies can move activities in and out more rapidly than before. But the basis for those decisions has broadened. Cost still matters, but so do latency, optionality, resilience, data ownership, customer experience, and strategic control.

3. AI Changes the Firm From The Inside

AI introduces a deeper shift because it changes coordination inside the company itself.

In traditional organizations, expanding scope usually meant adding people. More products, regions, customers, workflows, and compliance obligations required more managers, meetings, approvals, reporting lines, and oversight. Growth created coordination overhead. The larger the organization became, the more energy it spent managing itself.

AI can weaken that historical link between scope and headcount.

An AI system can route information, summarize exceptions, monitor performance, draft communications, trigger workflows, reconcile data, compare options, and execute tasks without adding another person to the coordination chain. The immediate productivity gain is useful, but the deeper organizational gain comes when AI removes the need for handoffs among people.

This distinction matters. If AI merely helps employees move faster, the organization remains fundamentally human-coordinated. The meetings, approvals, escalations, and verification loops still exist. But if AI can be trusted to complete certain tasks without human review at every step, the company can remove portions of the human dependency graph altogether.

That is when the economics change.

The firm can become broader without becoming proportionally larger. It can take on more activities, serve more customers, operate more workflows, and integrate more capabilities while adding fewer people than an industrial-era organization would have required. In some cases, the company may expand its true operational scope while reducing headcount.

This is why employment will become a less reliable measure of firm size. A company may look smaller by headcount while becoming larger by activity, reach, data, and control. The relevant question is not how many people the company employs, but how much trusted work it can coordinate.

The key empirical signal of the AI-native firm will be scope/headcount decoupling: firms that grow revenue, product scope, customer reach, operational complexity, and control while employment stays flat or falls. That is the measurable signature of a company using trusted AI to coordinate work that previously required layers of people.

4. The Trust Gap Is The Strategic Divider

AI does not have one uniform effect on company boundaries. It can strengthen the case for outsourcing or strengthen the case for internalization. The difference is trust.

The crucial question is whether a company trusts AI more inside the firm or outside it.

Internal AI may have important advantages. It can be trained or tuned on proprietary data. It can operate inside the company's security perimeter. It can be observed, audited, governed, and improved by people who understand the business context. It can be tied directly to internal workflows, policies, customer histories, risk rules, and operating constraints.

External AI may also have advantages. It may be cheaper, more advanced, more standardized, and more rapidly improved by a specialist provider. For routine, well-specified, low-risk tasks, the market may deliver better AI capability than any individual company can build.

The balance between these two conditions creates a Trust Gap.

When internal AI is more trusted, companies will tend to bring more activity inside. Proprietary data, auditability, security, and governance make internal systems more reliable for the company's distinctive work. In this world, AI does not shrink the firm. It can make the firm more powerful at larger scope because the company can coordinate more activity without reproducing the old bureaucratic burden.

When external AI is more trusted, companies will tend to outsource more. Standardized tasks that are easy to specify, measure, and verify will migrate to AI-enabled markets. In this world, the firm becomes thinner, retaining the activities where it has distinctive judgment, data, brand, or accountability obligations.

When internal and external AI are equally trusted, the boundary becomes less about operating efficiency and more about fixed institutional factors: legal structure, licensing, capital, brand, governance, and customer relationships. The company may become a lighter institutional shell around a network of AI-enabled services.

This is the central strategic implication: AI adoption alone does not predict whether firms get larger or smaller. The direction depends on where trusted agency resides.

Digital platforms flatten market coordination costs. AI flattens both market and internal coordination costs. The direction of firm scope therefore depends on the Trust Gap: whether internal or external AI can be trusted more. As technical coordination gets cheaper, accountability costs become first-order. The firm persists as the legal and institutional container for trusted agency.

5. Human-In-The-Loop Is Not A Detail. It Is The Business Model.

Many companies talk about keeping humans in the loop as if it were a governance preference. Economically, it is much more than that.

A human-in-the-loop process preserves the old coordination structure. The AI may draft, recommend, flag, or summarize, but humans still review and approve. That may be necessary in high-risk settings, but it limits the organizational transformation. The company still carries the cost and delay of human verification.

A human-out-of-the-loop process is different. It means the company is willing to let the system act within defined boundaries without routine human approval. That requires confidence in data quality, model behavior, monitoring, escalation rules, compliance controls, and accountability. It also requires clarity about what happens when the system is wrong.

This is where many AI strategies will succeed or fail.

The largest gains will not come from sprinkling AI assistants across existing workflows while leaving the structure untouched. They will come from redesigning workflows around trusted delegation. That requires executives to classify tasks by the level of autonomy they can safely support.

Some tasks should remain human-led because they require judgment, ethics, empathy, negotiation, creativity, or accountability. Some should be AI-assisted because the cost of error is meaningful but manageable. Some can become AI-executed because they are bounded, observable, reversible, and governed. Some should be outsourced to market AI providers because they are standardized and nonstrategic. And some should be internalized because they depend on proprietary context or carry unacceptable risk if delegated externally.

The point is not to remove humans everywhere. The point is to stop using human review as a default substitute for trust architecture.

6. Accountability Is The Reason Firms Do Not Disappear

As AI makes coordination cheaper, it is tempting to imagine a future of tiny companies orchestrating vast networks of autonomous agents. Some of that will happen. But it misses a hard institutional fact: society still needs someone to be responsible.

AI agents cannot hold a banking license. They cannot be sued in the ordinary commercial sense. They cannot sign a fiduciary duty, appear before a regulator, indemnify a customer, repair a damaged reputation, or absorb liability after a failure. Legal and commercial accountability attaches to people and institutions.

That is why the firm remains essential.

The firm is the structure through which society assigns responsibility for economic action. It holds contracts. It owns assets. It carries insurance. It stands behind warranties. It hires officers. It maintains compliance programs. It answers to regulators, courts, customers, investors, and counterparties.

In a world of AI agents, this accountability function becomes more, not less, important. When technical coordination is expensive, the firm exists partly to reduce coordination cost. When technical coordination becomes cheap, the firm exists increasingly to concentrate trust, liability, governance, and recourse.

This changes the executive lens. The make-or-buy decision must include not only operating cost, but also accountability cost. Leaders must ask:

- Who is responsible if the AI system fails?
- Can the vendor's performance be audited deeply enough?
- Is the task regulated or license-dependent?
- Would a failure harm customers, markets, patients, citizens, or critical infrastructure?
- Does outsourcing create a chain of dependencies no one can fully observe?
- Would the company still be blamed even if the external provider caused the failure?

The answer to these questions may determine the boundary of the firm more than the price of the technology.

This is the Fiduciary Wedge: the gap between what technology can coordinate and what institutions must still answer for. The term is broader than legal fiduciary duty in the strict sense. It also includes regulatory exposure, reputational liability, and the political cost of systemic failure. The wedge is not a fixed cost. In high-stakes systems, it compounds with scale.

7. When Outsourcing Becomes More Dangerous At Scale

The accountability problem is not constant. It grows with the scale and interconnectedness of the activity.

For a low-risk task, an external AI failure may be annoying but bounded. A marketing asset is off brand. A support answer needs correction. A back-office reconciliation requires cleanup. These failures matter, but they do not threaten the institution.

In high-stakes systems, the risk is different. A bad medical recommendation, a flawed credit decision, a trading error, a transportation failure, a cybersecurity breach, or a defense-related mistake can cascade. The problem is not just one vendor error. It is the interaction of many automated decisions across a complex network.

At small scale, outsourcing may still be sensible. A company can experiment, learn, and limit exposure. At large scale, the same outsourcing model can become dangerous because the company no longer has sufficient visibility or control over systemic risk.

This is where the classical logic of the firm can reverse.

In the industrial era, large firms often outsourced because internal bureaucracy became too heavy. In the AI era, large firms in high-liability sectors may internalize because external dependency becomes too risky. They will bring more capabilities inside not to save money, but to contain responsibility.

This will matter most in sectors where failure is regulated, systemic, or morally charged: finance, healthcare, energy, transportation, defense, insurance, public infrastructure, and sensitive data services. In these industries, the company cannot treat AI coordination as a simple procurement decision. It must decide which capabilities are too consequential to leave outside its own trust perimeter.

8. Three Emerging Organizational Patterns

The interaction of AI capability, trust, and accountability points toward three broad organizational patterns.

Pattern One: The Elastic Enterprise

In low-risk, digitally mature sectors, companies will move activities in and out of the firm more fluidly. They will rely heavily on platforms, vendors, contractors, and AI-enabled services. Internal teams will focus on brand, customer insight, product direction, partnerships, and high-leverage judgment.

These firms will be optimized for flexibility. Their boundaries will shift as market conditions change. They will avoid owning capabilities when external markets are faster, cheaper, and good enough.

Pattern Two: The Trusted-Core Enterprise

In sectors where proprietary data, internal context, and governance matter, firms will build a trusted AI core. They may outsource generic tools, but they will internalize the systems that learn from their unique data and execute their distinctive processes.

These companies may grow scope faster than headcount. They will look lean from the outside but become increasingly capable internally. Their competitive advantage will come from combining proprietary data, AI execution, workflow redesign, and institutional trust.

Pattern Three: The Accountability Enterprise

In high-liability sectors, firms will internalize critical AI capabilities because responsibility cannot be safely delegated. They will still use vendors and platforms, but the most consequential decision systems will sit inside the institution's own governance perimeter.

These firms will not necessarily be bureaucratic in the old sense. AI may reduce human coordination costs dramatically. But they will be institutionally heavy where licenses, auditability, resilience, and liability demand it.

Each pattern can be strategically valid. The mistake is assuming that one model will dominate every industry.

9. What Leaders Should Do Now

The next generation of organizational advantage will come from making sharper boundary decisions. Leaders should begin with a practical map of the company's activities.

Start by separating activities into four categories.

First are activities that are standardized, measurable, low-risk, and widely available from external providers. These are candidates for outsourcing to AI-enabled markets.

Second are activities that benefit from AI but still require human judgment, relationship management, ethics, or creative direction. These should be redesigned as AI-assisted workflows rather than fully automated ones.

Third are activities where the company has a proprietary data advantage or unique operating context. These are candidates for internal AI systems because the company may be able to trust its own AI more than market alternatives.

Fourth are activities where failure creates legal, regulatory, reputational, or systemic exposure. These require the strongest governance. In many cases, they belong inside the firm even if external providers appear cheaper.

Then ask six strategic questions:

- Where do we have data or context that outside providers cannot match?
- Which workflows still depend on human handoffs only because we have not built trust into the system?
- Which AI decisions can be made observable, auditable, reversible, and bounded?
- Which vendor dependencies would become dangerous if scaled across the enterprise?
- Which activities would customers, regulators, or courts still hold us responsible for, regardless of who performed them?
- Where could we expand scope without expanding headcount if AI were trusted to execute?

These questions move AI strategy out of the productivity sandbox and into corporate strategy. They force the company to decide what it must be able to do, what it can safely buy, and what it must be accountable for.

What To Watch

The strongest evidence for this thesis will not be headcount reduction alone. Watch for companies that grow revenue, product scope, customer reach, data control, and operational complexity while headcount stays flat or falls. That is scope/headcount decoupling. It will appear first where proprietary data, workflow redesign, and internal trust architecture allow AI to execute bounded work without routing every decision through human coordination layers.

Also watch for the opposite pattern in high-liability sectors: large firms reintegrating AI capabilities that vendors can technically provide but cannot credibly absorb responsibility for. In these cases, internalization will be driven less by efficiency and more by containment of liability, auditability, and institutional trust.

10. The New Logic Of The Firm

The old firm was built to solve a human coordination problem. Markets were costly to use, contracts were incomplete, information was slow, and people had limited capacity to process complexity. Hierarchy helped managers direct work without renegotiating every task through the market.

The digital firm was built around platforms, data flows, and modular interfaces. It could coordinate more activity across boundaries and rely on external ecosystems without owning every asset.

The AI-enabled firm will be built around trusted agency. Its most important asset will not be headcount alone, or even software alone. It will be the ability to decide where autonomous systems can act, where humans must judge, where markets can be trusted, and where the institution itself must stand behind the outcome.

This does not mean firms disappear. It means their rationale becomes more precise.

Some companies will become thinner as external AI markets absorb standardized work. Some will become broader as internal AI lets them expand scope without recreating bureaucracy. Some will become more integrated because accountability risk makes outsourcing unacceptable at scale.

The winners will not be the companies that adopt the most AI tools. They will be the companies that understand the new boundary problem.

What should we own? What should we orchestrate? What should we trust? What must we be responsible for?

Those are now the defining questions of strategy.

The firm of the AI era is not merely an employer, a production system, or a contracting hub. It is the institution through which trusted agency is assembled, governed, and made accountable.

Note For Economists

In formal terms, this argument modifies the Coasean boundary condition by allowing coordination costs to vary by technological regime. Digital platforms reduce market transaction costs. AI changes both internal and external coordination costs, with the direction of the boundary determined by relative trust in internal versus market AI. Accountability costs then become convex in high-liability or systemic environments, creating conditions under which AI can make firms broader in scope even as they become leaner in employment. This is not the death of Coase. It is Coase under machine coordination and institutional accountability.