Will AI Make Firms Bigger or Smaller?

Man Group / Man Institute · 2026-08-11 · English original

The answer is both. Which way an industry tips may reshape M&A activity before it shows up in the productivity statistics.

Key takeaways

Boardrooms debating AI strategy keep running into the same argument. One side wants to outsource more of what the firm currently does in-house, arguing that AI monitoring tools make suppliers easier to manage. The other side wants to bring more activities inside the firm, with the view that AI coordination tools make bigger organisations easier to run. Both sides have evidence, and in our view neither is wrong.

How AI reshapes firm boundaries (i.e. what it does in-house and what it buys from outside) is an active question in current economic research. Our contribution is a framework we have developed at Man Group for thinking through the question, and we should be clear at the outset that it is analysis rather than an empirical study. It builds on the transaction cost economics of Ronald Coase and Oliver Williamson set out in the next section and extends it to ask what AI does to those costs. The full formal model,1 with the underlying mathematics, is set out in a companion working paper. This article keeps to the reasoning behind it and what it means for portfolios.

Two parts of the framework are our own. The first is that AI acts on the firm's boundary in two opposing directions at once, making contract enforcement cheaper, which favours outsourcing, and internal coordination cheaper, which favours scale. The second is the data feedback loop, in which broader integration produces richer proprietary data, sharper AI and lower costs again, which strengthens the case for integrating further. The framework is theoretical. It reasons from characteristics that are measurable today to what we expect to see over the next few years, rather than reporting results we have already estimated.

So based on that, the outsourcing camp is right; AI monitoring tools can make supplier quality verifiable in real time, reducing the cost of market transactions. That said, the acquisition camp is right too, as AI coordination tools can compress the communication overhead inside large organisations. Both forces just apply to different industries, and we can already observe the characteristics that determine which way a given industry tips today.

Two forces, opposite directions

Companies do not have to exist. Anything a firm does in-house could in principle be bought from an outside supplier, so why draw a boundary and call it a company at all? Coase asked exactly that in 19372, and his answer still frames the debate. Firms exist because using the market is not free: finding suppliers, agreeing terms and policing the deal all cost money, what economists call transaction costs. When those costs climb high enough, it is cheaper to bring the work inside. Oliver Williamson, who won the 2009 Nobel prize (shared with Elinor Ostrom) for this work3, as Coase did in 1991, later sharpened the point.

He found that the costly part is less striking the deal than enforcing it when a dispute arises, especially where one side has sunk money into assets built for that specific relationship and worth little outside it. AI now acts on both costs at once, which is why a question posed nearly ninety years ago is still alive in boardrooms today.

Figure 1: How AI reshapes firm boundaries, by industry type

Industry classification by relationship-specific investment and coordination complexity

Source: Man Group. Framework from Bond (2026), “Artificial Intelligence and the Boundary of the Firm: A Coasean Analysis,” Man Group working paper.

Coase and Williamson wrote long before the advent of AI, but we believe their cost framework is the right lens for what it does. Apply it, and AI turns out to act on each cost in opposite directions.

How AI can make firms smaller

On the market side, AI can make contracts easier to enforce. Digital records, whether automated audit trails, real-time quality monitoring, or anomaly detection, produce evidence that holds up when disputes arise. The mechanism here is evidentiary rather than a coordination one, as it does not require AI systems to interact across firm boundaries, only that digital records make contractual performance provable to courts. When you can prove what happened, you don't need to own the supplier to protect yourself. So the case for keeping things in-house weakens.

This force is strongest in industries where the deliverable is bespoke or relationship-specific and contracts are inevitably incomplete, such as specialised engineering, custom manufacturing and commercial real estate brokerage. Disputes in these industries have historically been expensive because the evidence of what was actually delivered was incomplete or contested. AI-generated records make that evidence harder to dispute.

There is an important caveat. This depends on courts actually giving AI-generated records evidentiary weight in disputes, which is governed by evolving rules on electronic and algorithmic evidence (in the US, FRE 901–902 and the Daubert standard). Electronic-transaction regimes such as the EU's eIDAS or Singapore's Electronic Transactions Act establish that digital records are valid, which helps, but the binding question is whether courts will admit AI-generated operational evidence. Where those standards are maturing the channel is already active; elsewhere, “black box” AI outputs face the same scepticism as any unverifiable claim. The legal environment is a leading indicator of how fast this channel takes effect.

How AI can make firms bigger

On the internal side, AI makes large organisations cheaper to run. The fundamental constraint on firm size has always been coordination, i.e. the cost of getting information to the right person at the right time across layers of management. AI compresses that cost through automated workflows, decision-support tools and AI-driven resource allocation.

This force is strongest in industries with high coordination complexity such as logistics networks, multi-stage manufacturing or large-scale retail. A rough proxy is whether the CEO's biggest operational challenge is getting divisions to talk to each other. Where it is, coordination complexity is high.

A second caveat is that AI tends to absorb the routine decisions first, so the ones still reaching human managers are the harder cases. In industries that depend on irreducible expert judgment, the coordination gains may be partly offset by rising decision complexity, leaving an ambiguous net effect.

The deciding factor is data

The static picture, that enforcement gets easier, so some firms tend to shrink, while cheaper coordination tends to help others grow, is useful but incomplete. The more interesting result centres around data.

A firm that brings more related activities in-house accumulates more proprietary operational data. A logistics company managing a broader network generates more route-optimisation data. A manufacturer controlling more production stages generates more process data. Data you can buy gets commoditised; data only your own operations can generate does not. That data makes a firm's AI systems more effective, which further lowers its coordination costs, which typically makes even broader integration attractive.

Amazon's steady march into logistics, from warehousing, to last-mile delivery and air freight, is the kind of pattern we have in mind, though we wouldn't call it proof. Plain economies of scale and network effects explain much of it too, and most of it predates today's AI. What's new is the reason to expect more of it, each step generates operational data that makes the next one cheaper to run.

The same logic plays out in contested industries, those where both contract-enforcement gains and coordination gains are substantial. Healthcare and financial services are cases in point. A hospital system that integrates across specialties accumulates richer operational data from patient flows and resource utilisation to cross-departmental scheduling. That likely lowers its coordination costs, making further integration attractive. A rival without that data breadth tends to outsource to specialists instead. Two firms in the same industry end up with opposite strategies, and the difference is which one first crossed the data threshold and set the feedback loop turning. Contested industries therefore produce the widest within-sector dispersion because the gap between data-rich integrators and data-poor specialists is largest where both forces are strong.

Why the loop favours scale, not enforcement

This is a positive feedback loop, but it applies asymmetrically. Evidentiary data (sensor logs, audit trails) is task specific. Proof that one supplier delivered adequate steel does not help verify another supplier's software. Coordination data works the other way and is cross-task, so route-optimisation data from one logistics leg improves scheduling across the entire network. The feedback therefore compounds on the coordination side while the contracting side barely moves.

In industries where coordination costs are the binding constraint, it produces two self-reinforcing outcomes:

Figure 2: The Data Flywheel - data feedback loop above and below the critical data threshold

Source: Man Group. Framework from Bond (2026), "Artificial Intelligence and the Boundary of the Firm: A Coasean Analysis," Man Group working paper.

Winner-take-most, and its limits

Where data advantages compound strongly enough, the likely result is winner-take-most with the firms that integrate broadly ending up both bigger and lower-cost, provided that expanding doesn't dilute what the firm is actually good at. On that condition, for standardised products and services, staying small turns out to be the more expensive model rather than simply an alternative one.

One point is important to clarify. We are talking about concentration among the companies that use AI, meaning data-rich operators in logistics, manufacturing or healthcare, not the AI providers themselves, where lookalike models may compete for margins instead of handing any one of them the market. The point concerns ordinary companies building data advantages in their own industries, rather than the question of who wins among the model builders.

What widens the gap, and what bounds it

Synthetic data, which means artificial data a firm's own AI generates to train on when real-world data is scarce, widens that gap rather than closing it. Data-rich firms can augment their proprietary datasets with AI-generated synthetic data, validated by their own operations, and the effect compounds over time.

Data-poor firms get the opposite, because training on recycled synthetic outputs risks degrading model quality rather than improving it, a dynamic researchers call model collapse (Shumailov et al., 2024).4 Rising cloud exit costs make the low-integration trap worse still, because firms whose operational data sit on a provider's infrastructure face a choice between losing access to that data or absorbing prohibitive migration costs.

Data advantages fade fast between unrelated activities. A logistics firm gains far more from adding adjacent route segments than from moving into unrelated businesses, and AI model performance eventually saturates in any case. The winner-take-most dynamic is therefore self-limiting. Even in consolidating industries, the outcome is a dominant firm within a related set of activities, rather than a monopoly over everything.

A practical map

Our framework classifies industries along two dimensions. The first is the relationship-specificity of key investments, which determines how much AI helps contract enforcement. The second is the coordination-intensity of the business, which determines how much AI helps internal management.

Where an industry sits on the two dimensions decides which of four futures it could face, and the differences between them are sharp.

Consolidation (high coordination, low specificity): AI coordination gains dominate and data feedback reinforces scale, so logistics, integrated manufacturing and large-scale retail tend to concentrate.

Fragmentation (low coordination, high specificity): AI-generated evidence makes outsourcing cheaper, so specialised engineering, custom manufacturing and commercial real estate brokerage tend to fragment.

Contested (both high): the outcome depends on which firm builds the data advantage first, which produces the widest dispersion within a sector. Healthcare and financial services sit here. This quadrant gives a spread of outcomes rather than a directional call, and it is also where data-sharing regulation would bite hardest.

Status quo (both low): AI has limited organisational impact, as in artisanal production and personal services.

What it means for portfolios

We think M&A will increasingly be driven by data more than scale. In coordination-heavy sectors, the framework points to a distinct reason to acquire. Firms will buy targets for the operational data they bring, data that compounds the buyer's AI advantage, and no longer for conventional scale or cost savings. Deal rationales in logistics, manufacturing and retail that turn on data rather than footprint or headcount are the early signature of this shift.

Concentration risk is sector-specific. The data feedback loop predicts increasing concentration in the high-coordination quadrant like logistics, manufacturing and retail, where firms with early data advantages have a self-reinforcing moat. In relationship-heavy sectors (specialised engineering, custom manufacturing, commercial real estate brokerage), the opposite happens where incumbents lose their integration advantage as contracts become easier to enforce, and industries fragment toward smaller, more specialised firms.

Data advantages plateau eventually, but the gap between the leader and the rest tends to widen first. In the meantime, the firms positioned to benefit are those that have already integrated broadly and are investing in proprietary data infrastructure, the operational businesses rather than the specialist AI vendors. Two observable signs that a firm is crossing the data threshold are the breadth of its supply-chain integration and the proprietary data assets it discloses in regulatory filings.

The legal environment is a leading indicator. As noted above, the pace of the fragmentation shift depends on courts accepting digital evidence. That cross-border variation is an underappreciated source of sector dispersion: the same industry may consolidate in one jurisdiction and fragment in another.

How the framework reads across a sector, and what markets may already reflect. The clearest signal is a relative one. Within a sector, the framework points to a widening gap between broadly integrated operators that disclose proprietary-data infrastructure and data-poor specialists, and that gap should become clearer as integration metrics turn observable. Two cautions temper this. The first is crowding. Favouring the data-rich integrator overlaps heavily with the consensus mega-cap AI-winner trade already in prices, so any edge lies more in distinguishing winners from losers inside a sector than in the sector call itself. The second concerns capital spending. AI-related capex as a share of revenue is now near-universal among large caps and tells you little on its own, so the more revealing questions are how firm-specific the data is and whether the spend is aimed at integration. Above all, this remains a theoretical prior awaiting the tests set out in the working paper, better treated as a direction for research than a validated signal.

Our bottom line

Whether AI makes firms bigger or smaller depends first on what "bigger" means. Here it means broader scope, more activities inside the firm's boundary, and not necessarily more employees. On that definition the answer is both, and the characteristics that decide which way an industry tips are measurable today: relationship-specific investment, coordination complexity, data endowments and legal infrastructure. A firm can consolidate its supply chain or expand across related specialties while shrinking its headcount. So the measure to watch is integration, both vertical and horizontal, ahead of payroll.

The aggregate productivity debate matters but moves slowly. The firm-boundary question moves faster. The data feedback loop, with first movers pulling away from late entrants, is the dynamic most likely to generate investable dispersion across sectors before the macro data catch up. Signs should start to show in vertical integration measures such as M&A deal flow, census segment data and make-versus-buy ratios, though the timing is hard to call precisely. Three developments would undermine the thesis. Courts could reject AI-generated evidence, which would stall fragmentation. Antitrust action could target data-driven integration, which would stall consolidation. Or data-sharing rules, such as open-banking-style mandates and the EU's Data Act, could hand laggards the data they lack, letting them catch up and close the gap between leaders and the rest that the whole thesis depends on.

The framework above offers a way to answer which industries consolidate, which fragment and how fast, before the market does.

Note: Additional reading - Companion macro assessment: Gregory Bond, “The Productivity Paradox: When Will AI Deliver?,” Man Group, February 2026. On the AI infrastructure and valuation cycle, see also “The AI Bubble: Hidden Risks and Opportunities,” Man Group/Oxford Man Institute.

1. Bond, G. (2026). “Artificial Intelligence and the Boundary of the Firm: A Coasean Analysis." Man Group working paper. 2. Coase, R.H. (1937). “The Nature of the Firm.” Economica, 4(16), 386–405 3. Williamson, O.E. (1985). The Economic Institutions of Capitalism. New York: Free Press. 4. Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). "AI Models Collapse When Trained on Recursively Generated Data." Nature, 631, 755–759

For further clarification on the terms which appear here, please visit our Glossary page.
