AI and the Nature of the Firm

Alec Litowitz
Qstar Capital

Nick Polson
Booth School of Business
University of Chicago

Vadim Sokolov
Department of Systems Engineering
and Operations Research
George Mason University

First Draft: April 2026
This Draft: April 16, 2026

Abstract

Artificial intelligence is a horizontal general-purpose technology that restructures the cost functions of every knowledge-intensive industry simultaneously. We develop a unified Coasian framework for analyzing this transformation, drawing on three of Coase's contributions: the theory of the firm, the Coase conjecture on durable-goods monopoly, and the Coase theorem on efficiency under low transaction costs. Our central argument is that AI operates as a direct solvent on information asymmetry rents, the foundational friction that justifies the existence of large knowledge-economy firms, professional service intermediaries, and high-margin software businesses. Beyond reducing transaction costs, AI automates the directing function, the exercise of agency over resource allocation, that Coase identified as the firm's core rationale. Using Coase's theory of the firm, we show that AI reduces external transaction costs faster than internal coordination costs, predicting a contraction in firm size coupled with an expansion in market-mediated exchange. Through the Coase conjecture, we argue that AI capabilities function as a durable good whose producer competes against its own future price reductions, driving knowledge-processing rents toward zero. Through the Coase theorem, we argue that as transaction costs fall, the initial allocation of knowledge assets becomes less determinative of efficient outcomes, enabling reallocation across firm boundaries. Early empirical evidence from labor market studies supports the theoretical predictions: occupations with high observed AI exposure show weaker projected employment growth, though aggregate displacement remains modest. The framework predicts that the firms which survive and concentrate value will be those with durable advantages in distribution, proprietary data, trust, regulatory position, and judgment, not information processing.

1 Introduction

What happens to the structure of the knowledge economy when the cost of processing, verifying, and transmitting information collapses? This paper argues that artificial intelligence is best understood not as a product innovation but as a horizontal general-purpose technology, analogous to electricity, that dissolves the informational frictions sustaining the current architecture of firms, professional services, and software markets.

The analytical foundation is Ronald Coase's theory of the firm [Coase, 1937]. Coase showed that firm boundaries are determined by the relative cost of internal coordination versus external market transactions. A firm expands when bringing an activity inside is cheaper than contracting for it in the market; it contracts when the reverse holds. But Coase's insight runs deeper than transaction costs. Inside the firm, the entrepreneur replaces the price mechanism as the director of resources. The firm exists not only because market transactions are costly, but because the act of directing, of exercising agency over resource allocation, is sometimes more efficient than relying on decentralized price signals. The firm is a locus of agency.

AI challenges this structure on both dimensions. It reduces transaction costs, making market-mediated exchange cheaper. But it also automates the directing function itself. When an AI system evaluates candidates, allocates budgets, routes decisions, or drafts strategies, it performs the entrepreneurial function that Coase identified as the firm's core rationale. The implication is that AI does not merely shift the Coasian boundary; it competes with the human agency that constitutes the firm. Workers become less like permanent members of a hierarchy and more like callable functions: modular capabilities invoked when their specific judgment, creativity, or embodied knowledge is required, and released afterward. As Litowitz et al. [2026b] argue, the binding constraint in the AI economy is not how many questions can be answered but which questions are worth asking, a problem of agency and direction that computation alone cannot solve.

The knowledge economy amplified Coasian logic: when the relevant assets are expertise, institutional knowledge, and verified judgment, they are sticky, hard to specify contractually, and subject to severe asymmetric information problems. Knowledge-intensive firms grew large not primarily to achieve scale economies in production but to internalize the transaction costs associated with these information assets. Many entire knowledge industries exist as frictions: the legal profession intermediates information asymmetry between parties; management consulting exists because firms cannot cheaply verify external strategic advice; financial advisory intermediates the gap between what investors know and what they need to know. These industries do not produce goods; they reduce transaction costs. When AI drives those costs toward zero, the industries built on them lose their economic rationale.

AI dissolves these frictions asymmetrically. External transaction costs for knowledge work fall dramatically; internal coordination costs fall less uniformly, because organizational politics and status competition are not information problems and AI cannot dissolve them.

We extend the Coasian analysis in two directions. First, we connect it to the Coase conjecture [Coase, 1972], which shows that a durable-goods monopolist competing against its own future pricing decisions is driven to price at marginal cost. We argue that AI capabilities exhibit the economic properties of a durable good, and that providers of AI-powered knowledge services face the same self-competition dynamic: each generation of capability makes prior pricing unsustainable, driving information-processing rents toward zero. Second, we invoke the Coase theorem [Coase, 1960] to argue that as AI drives transaction costs in knowledge markets toward zero, the initial allocation of knowledge assets, including institutional expertise, proprietary workflows, and organizational memory, becomes less determinative of efficient outcomes. Knowledge can be reorganized across firm boundaries at low cost, and the question of who "owns" institutional knowledge becomes less economically meaningful.

Early empirical evidence is consistent with these predictions. Massenkoff and McCrory [2026] introduce a measure of "observed exposure" combining theoretical AI capabilities with actual usage data, finding that occupations with higher observed exposure show weaker projected employment growth through 2034. At the same time, aggregate unemployment effects remain modest, consistent with the Coasian prediction that the transition involves reorganization rather than simple displacement: tasks migrate across firm boundaries rather than disappearing.

The paper makes three contributions. First, it provides a unified Coasian framework, drawing on the theory of the firm, the Coase conjecture, and the Coase theorem, for analyzing how AI restructures the knowledge economy. Second, it identifies the specific mechanism through which AI creates value and destroys rents: the dissolution of information asymmetry, the foundational friction that justifies knowledge-economy firm size, professional intermediation, and software switching costs. Third, it characterizes the "residual firm," the organizational form that emerges when information-processing rents collapse, and identifies the durable sources of value that survive: distribution, proprietary data, trust, regulatory position, and human judgment.

Section 2 develops the horizontal technology framing. Section 3 presents the Coasian theory of firm boundaries. Section 4 develops the agency dimension: AI as a directing technology, humans as callable functions, and the recursive creation of AI agents. Sections 5 and 6 extend the analysis through the Coase conjecture and the Coase theorem. Section 7 reviews early empirical evidence. Section 8 characterizes the residual firm, and Section 9 discusses implications and limitations.

2 AI as a Horizontal Technology

The distinction between vertical and horizontal technologies is fundamental to understanding AI's economic impact. A vertical technology is a superior product that creates a new market or displaces an existing one: the smartphone replaced the feature phone, the digital camera replaced film. A horizontal technology does not merely substitute for a prior input but reorganizes every production process that depends on it.

The canonical historical example is electrification [David, 1990]. Before electric power, factories were built vertically around a central drive shaft connected to a water wheel or steam engine. Power had to be physically proximate to its source, and the entire layout of production was dictated by this constraint. Electrification dissolved the constraint entirely: factories reorganized horizontally around workflow rather than power source.

The resulting productivity gains came not from cheaper energy per se but from the architectural reorganization that cheap, ubiquitous energy made possible. David [1990] documents that the full productivity gains of electrification took over forty years to materialize, because they required complementary reorganization of production processes, management practices, and worker skills.

AI exhibits the same horizontal character, but the friction it dissolves is informational rather than physical. The cost of translating, verifying, synthesizing, and routing knowledge across organizational boundaries has historically been the binding constraint on the architecture of knowledge work. This friction sustains the high-margin professional services complex (law, consulting, financial advisory, compliance) and the large internal knowledge hierarchies of modern corporations.

Bresnahan and Trajtenberg [1995] formalize the concept of a general-purpose technology (GPT) as one that (i) is pervasive, used across many sectors; (ii) improves over time, creating sustained incentives for complementary innovation; and (iii) generates innovation spillovers, making it easier to invent in application sectors. AI satisfies all three criteria. It is deployed across virtually every knowledge-intensive industry. Its capabilities improve rapidly, with each generation expanding the frontier of automatable tasks. And it reduces the cost of experimentation and prototyping across application domains, from drug discovery to software development to legal analysis.

The GPT framework predicts a characteristic pattern: an initial period of disruption and uneven adoption, followed by broad productivity gains once complementary innovations (organizational restructuring, new business models, updated regulatory frameworks) catch up. This pattern, documented by David [1990] for electricity and by Brynjolfsson and Hitt [2000] for information technology, is consistent with the early evidence on AI adoption, which shows rapid capability gains but modest aggregate productivity effects thus far.

3 The Coasian Theory of the Firm Under AI

This section develops the central theoretical argument: AI shifts the Coasian boundary of the firm inward by reducing external transaction costs faster than internal coordination costs.

3.1 Transaction Costs and Firm Boundaries

Coase [1937] posed a deceptively simple question: if markets are efficient at allocating resources, why do firms exist? His answer was transaction costs. Using the price mechanism is costly: discovering relevant prices, negotiating contracts, specifying terms, and enforcing agreements all consume resources. When these costs are high, it is cheaper to organize production within a firm under hierarchical authority. The boundary of the firm is set where the marginal cost of internal coordination equals the marginal cost of market transaction:

$$MC_{\text{internal}} = MC_{\text{market}}. \tag{1}$$

Williamson [1975] extended this framework by identifying three key frictions that drive the internalization decision: asset specificity (investments that lose value outside a particular relationship), opportunism (the risk that contracting partners act in self-interest with guile), and bounded rationality (the cognitive limits on managers' ability to process all relevant information). These frictions are particularly acute for knowledge work, where the relevant assets are intangible, hard to specify contractually, and subject to severe information asymmetries.

The knowledge economy amplified Coasian logic in a specific direction. When the relevant production inputs are expertise, institutional knowledge, and verified judgment, they are sticky and hard to transfer. A law firm's value lies not in its physical plant but in the accumulated knowledge of its partners, their relationships with clients, and their understanding of institutional norms. This knowledge is tacit in the sense of Polanyi [1966]: "we know more than we can tell." The result is that knowledge-intensive firms grew large primarily to internalize the transaction costs associated with knowledge assets, not to achieve scale economies in production.

3.2 AI Shifts Both Cost Curves

AI shifts both legs of the Coasian trade-off simultaneously, but asymmetrically.

External transaction costs fall dramatically across multiple dimensions. Knowledge work that previously required a permanent employee hierarchy, because the knowledge was too tacit, too context-dependent, or too costly to verify externally, can now be contracted, executed, and verified at low cost through AI systems. Table 1 summarizes the shift across the principal categories of market frictions in the transaction cost economics literature [Coase, 1937, Williamson, 1985].

Table 1: Transaction cost and coordination friction shifts under AI

Friction	Pre-AI	Post-AI
Search and matching	Costly identification and vetting of specialists	Near-zero cost to locate and deploy specialized knowledge
Specification	Writing detailed specs for complex knowledge tasks	AI drafts specifications, interprets intent from partial instructions
Verification	Expert review of output quality	AI monitors quality cheaply and continuously
Contracting	Legal drafting, compliance checking	AI handles most drafting and verification
Bounded rationality	Hierarchies to simplify decision-making	AI expands individual information-processing capacity

Internal coordination costs also fall, but less uniformly. AI reduces informational frictions within firms, including synthesis, search, and translation across functional silos. However, it does not reduce political frictions. Organizational politics, status competition, and the dynamics of authority are not information problems; they are problems of preference aggregation and power, and AI cannot dissolve them. As Jensen and Meckling [1976] observed, the agency costs within firms arise from misaligned incentives and private information. AI can reduce the information asymmetry component of agency costs by enabling cheaper monitoring and verification, but it cannot eliminate incentive misalignment.

3.3 Predicted Firm Structure

The asymmetric reduction in transaction costs generates a clear directional prediction.

Proposition 1 (Firm contraction)**.** *If AI reduces external transaction costs for knowledge work faster than it reduces internal coordination costs, then the equilibrium firm size in knowledge-intensive industries decreases: the Coasian boundary in equation* (1) *shifts inward.*

This follows directly from Coase's framework. When external transaction costs fall relative to internal coordination costs, activities previously internalized within firms become cheaper to procure through market transactions. The firm contracts.

The prediction is not that firms disappear but that they shrink in headcount while potentially maintaining or increasing revenue and scope. The layer that thins is the middle: knowledge workers whose primary function was information translation, verification, and routing between organizational levels. What survives is high-level judgment at the top and direct execution at the edges.

A natural counter-argument applies the Jevons paradox to firm boundaries: if AI makes coordination cheaper in all directions, firms might expand scope and grow larger rather than smaller. This is logically possible, and for firms whose growth is currently constrained by information-processing bottlenecks, AI may indeed enable expansion. The asymmetry argument rests on the observation that internal coordination costs include a large political component (status competition, turf wars, organizational inertia) that AI does not reduce, while external transaction costs are primarily informational and therefore more susceptible to AI-driven compression. To the extent that political frictions dominate within large organizations, the asymmetry holds. In organizations where political frictions are minimal (small teams, founder-led firms), the prediction is weaker.

Firm boundaries should also become more dynamic. If Coasian boundaries are set by relative transaction costs, and AI drives external transaction costs for knowledge work toward zero, firms should become far more porous and reconfigurable, assembling teams and capabilities for specific projects and dissolving them afterward rather than maintaining permanent hierarchies. The film industry has long operated on this model: production companies assemble specialized talent (directors, cinematographers, editors, actors) for a single project and dissolve the team upon completion [DeFillippi and Arthur, 1998]. AI extends the viability of this project-based organizational form to a wider range of knowledge-intensive industries.

3.4 Agency Theory Under AI Monitoring

Jensen and Meckling [1976] formalized the principal-agent problem: agents possess private information and potentially misaligned incentives, generating agency costs that justify internalizing activity within firms. AI alters this calculus by reducing the information asymmetry component of agency costs.

If AI can monitor output quality cheaply and continuously, the information gap between principal and agent shrinks. Principals can contract more efficiently with external agents because verification costs collapse. The organizational implication is a shift from the Chandlerian hierarchy [Chandler, 1977] toward a dynamic nexus of AI-mediated contracts, closer to Alchian and Demsetz [1972]'s team production model but with AI as the residual monitor rather than an owner-manager.

This does not eliminate agency problems entirely. Incentive misalignment persists even under perfect monitoring: an agent may comply with observable metrics while undermining harder-to-measure objectives. The residual agency problem under AI is one of goal specification rather than information asymmetry, a distinction with significant implications for organizational design.

4 Agency, Direction, and the Callable Function

The transaction cost framework developed in the preceding section captures one dimension of AI's impact on firm structure: the relative cost of market versus hierarchy. But Coase's original argument contains a second, deeper dimension that transaction cost economics has largely subordinated: the question of agency. This section develops that dimension and its implications for the organization of work.

4.1 The Entrepreneur as Director

Coase's 1937 formulation is explicit. Inside the firm, "the distinguishing mark of the firm is the supersession of the price mechanism" [Coase, 1937]. The entrepreneur replaces the market as the allocator of resources. Where the price mechanism coordinates through decentralized signals, the entrepreneur coordinates through direction: deciding what to produce, how to produce it, and who should do what. The firm exists because this act of conscious direction is, under certain conditions, cheaper than relying on prices.

Simon [1951] formalized this insight. The employment contract is distinguished from a sales contract by the transfer of authority: the employee agrees to follow the employer's direction within a zone of acceptance, rather than contracting for a specific output. The firm is not merely a cost-minimizing structure but an authority relationship. What makes the firm a firm, as opposed to a web of market contracts, is that someone inside it exercises agency, making decisions under uncertainty about how resources should be deployed.

Hayek [1945] identified the knowledge constraint on this directing function. The central economic problem is not the allocation of given resources but the use of knowledge dispersed across millions of individuals. No single director, however capable, can possess the local, time-sensitive, contextual knowledge that the price mechanism aggregates automatically. The firm's directing function is therefore bounded: it works for activities where the relevant knowledge can be centralized, and it fails where it cannot. Firm boundaries reflect, in part, the boundaries of what a director can know.

4.2 AI as a Directing Technology

AI does not merely reduce the costs of directing. It performs the directing function itself. When an AI system triages customer inquiries, routes tasks to workers, evaluates vendor proposals, drafts hiring criteria, allocates advertising budgets, or recommends strategic priorities, it exercises the entrepreneurial function that Coase and Simon identified as the firm's constitutive feature. The system does not simply provide information to a human decision-maker; it makes the decision.

This is a qualitative shift. Transaction cost economics treats the directing function as fixed and asks how far it should extend. The question was always "should we make or buy?" with a human director making the call either way. AI introduces a third option: the directing itself can be automated. The make-or-buy decision can be made by a machine, and so can the decision about what to make, whom to buy from, and when to switch.

The implications cascade. If the firm's core function is direction under uncertainty, and AI can direct, then the Coasian rationale for the firm weakens on both legs simultaneously: market transactions become cheaper (the transaction cost channel) and the internal directing function becomes automatable (the agency channel). Workers in this regime are no longer permanent members of a hierarchy receiving ongoing direction from an entrepreneur. They become callable functions: modular capabilities with defined interfaces, invoked when their specific expertise is required and released when it is not. The firm becomes an orchestration layer that assembles human and artificial capabilities on demand, rather than a repository of permanent employees receiving continuous direction.

This organizational form is already visible in AI-native companies, where small teams use AI agents to perform functions that previously required departments: market research, legal review, financial modeling, customer support, content production. Each function is called as needed rather than maintained as a standing capability.

4.3 Agents Creating Agents

Good [1965] anticipated a further consequence. He defined an ultraintelligent machine as one that "can far surpass all the intellectual activities of any man however clever." Good's key observation was recursive: if such a machine could design machines better than itself, the result would be an "intelligence explosion" in which the first ultraintelligent machine is "the last invention that man need ever make."

The organizational analog is already emerging in a weaker form. AI agents can now design, deploy, and manage other AI agents. An orchestrating AI can decompose a complex task, spawn specialized sub-agents, evaluate their outputs, and synthesize results, performing the managerial function of task decomposition and quality control that was previously the exclusive province of human directors. The directing function is not merely automated but made recursive: agents create and direct other agents.

This recursion dissolves the principal-agent hierarchy that Jensen and Meckling [1976] analyzed. In the classical framework, a human principal delegates to human agents and bears monitoring costs. In the emerging framework, an AI principal delegates to AI agents, with human judgment invoked only at the boundaries: setting the objective, evaluating outcomes that require contextual or ethical judgment, and bearing accountability for the result. The human role shifts from continuous direction to intermittent oversight and goal specification.

Good's intelligence explosion remains speculative. But the organizational explosion, the rapid proliferation of AI-mediated task decomposition and delegation, is observable. The question is not whether machines will surpass human intelligence in some general sense, but whether they will absorb enough of the directing function to restructure how firms organize work. The evidence from Sections 3 and 7 suggests they already are.

4.4 What Agency Remains Human

If AI absorbs the directing function, what remains of human agency in economic organization? Three categories resist automation.

Goal specification. AI systems optimize objectives; they do not originate them. Deciding what a firm should pursue, what trade-offs to accept between profit, growth, risk, and social impact, requires a judgment that integrates values, politics, and vision. This is the Hayekian problem applied to ends rather than means: the question of which questions are worth asking [Litowitz et al., 2026b].

Accountability. AI can make decisions but cannot bear responsibility for them. Legal liability, fiduciary duty, and moral accountability require a human (or a human-controlled entity) that can be held answerable. The firm persists in part because someone must sign, someone must be sued, and someone must appear before a regulator. Agency in the legal sense remains irreducibly human.

Novel problem recognition. AI operates within the distribution of its training data. Recognizing that a genuinely new problem exists, one that the existing framework does not address, requires the capacity to notice what is missing rather than to process what is present. This is the frontier of the information ladder identified by Litowitz et al. [2026a]: codifiable knowledge approaches commodity pricing, but the identification of what is not yet known remains a human function.

These three residual domains of human agency, goal specification, accountability, and novel problem recognition, define the irreducible human core of the firm in the AI era. They are not information-processing tasks; they are exercises of judgment, responsibility, and creativity that current AI systems cannot perform. The residual firm, characterized in Section 8, is organized around these functions.

5 The Coase Conjecture and AI Economics

The Coase conjecture [Coase, 1972] addresses a different problem from the theory of the firm, but its logic maps directly onto the economics of AI-powered knowledge services. The conjecture concerns a durable-goods monopolist who cannot commit to future prices.

Because the good is durable (it does not depreciate), the monopolist competes against its own future pricing decisions. Rational consumers, anticipating future price reductions, defer purchases, and the monopolist is driven to price at marginal cost.

5.1 The Formal Structure

Consider a monopolist selling a durable good to consumers with heterogeneous valuations. Let consumers be indexed by their willingness to pay, with the highest-value consumer willing to pay y and the lowest willing to pay some value approaching zero. With zero marginal cost of production and a discount factor $\delta \in (0, 1)$ between periods, the monopolist faces the following problem.

If the monopolist charges y in period one, lower-valuation consumers wait. The monopolist must then lower the price in period two to capture additional demand. But the highest-valuation consumer, anticipating this, also waits. To make the high-value consumer indifferent between purchasing now and waiting, the monopolist must charge at most

$$p_1 = \delta x + (1 - \delta)y, \tag{2}$$

where x is the second-period price. As $\delta \to 1$ (consumers become patient), $p_1 \to x$, and the monopolist loses all pricing power. With n consumer types, the first-period price becomes $(1 - \delta)^n y$, which approaches zero as $\delta \to 1$.

The monopolist "competes against its own future self," and the competitive pressure drives the price to marginal cost despite the monopoly position.

5.2 AI Capabilities as a Durable Good

AI capabilities exhibit the economic properties that trigger the Coase conjecture. The knowledge-processing capability delivered by an AI system is durable: once a user has access to a model that can perform legal research, financial analysis, or code generation, that capability does not depreciate with use. Each new generation of AI capability is superior to the last, and the pace of improvement is rapid and publicly observable.

This creates a situation formally analogous to Coase's durable-goods monopolist. An AI provider selling knowledge-processing services today competes against its own future, more capable, and likely cheaper offerings. Rational buyers of AI-powered services, whether firms purchasing enterprise AI or individuals subscribing to AI assistants, anticipate that capabilities will improve and prices will fall. The incentive to defer adoption, or to demand lower current prices, is structurally identical to the durable-goods consumer's incentive to wait.

The implications extend beyond AI providers to every business whose margins depend on information-processing rents. If AI makes information processing a near-commodity input, then the firms that charge premium prices for information intermediation (consulting firms, financial advisors, legal services, SaaS companies encoding institutional knowledge) face the same Coasian self-competition dynamic. Each improvement in AI capability reduces the sustainable premium for human information processing. Litowitz et al.

[2026a] document this pattern across the AI value chain, showing that Hayek's knowledge problem remains unsolved for tacit knowledge even as codifiable knowledge approaches commodity pricing, and that Coasian firm structure is being bifurcated between small AI-augmented firms and large knowledge-fortress incumbents.

5.3 The Collapse of Software Terminal Value

The Coase conjecture connects directly to a parallel disruption in software economics. The traditional SaaS valuation model rested on three structural assumptions: high switching costs (software encodes institutional knowledge, creating lock-in), recurring revenue predictability (high retention justified growth-stage multiples), and expanding margins (software scales at near-zero marginal cost).

AI weakens all three. Switching costs collapse because AI can extract, replicate, and transfer the institutional knowledge encoded in enterprise software. The lock-in was always an information asymmetry rent; AI dissolves the asymmetry. Recurring revenue becomes contestable because competitive products can be built rapidly using AI-assisted development, shrinking barriers to entry. Margin assumptions break because AI-native entrants can enter at near-zero fixed cost, so price competition begins from day one rather than being deferred until after the incumbent has locked in its customer base.

The financial market implication is that the option value of future monopoly rents has collapsed. Markets historically priced high-growth software companies as if current growth rates predicted durable future pricing power. The Coase conjecture provides the theoretical explanation: when capability is durable and improving, the producer cannot credibly commit to maintaining current prices, and rational buyers discount future rents accordingly.

5.4 The Subscription Mitigation and Its Limits

The textbook response to the Coase conjecture is leasing rather than selling: if the monopolist retains ownership and charges a per-period rental fee, it avoids competing against its own past sales. AI providers have adopted precisely this strategy, overwhelmingly using subscription pricing rather than perpetual licenses. This is the correct commitment device, and it partially mitigates the conjecture's force.

Three features of the AI market limit the effectiveness of this mitigation. First, and most consequentially, the conjecture's logic applies not to AI subscriptions per se but to the information-processing rents that AI dissolves. A consulting firm cannot "lease" its expertise in the way that mitigates the conjecture; it can only charge what the market will bear, and AI steadily reduces that willingness to pay. The subscription model protects AI providers' revenue, but it does not protect the downstream firms whose margins depend on information intermediation. Second, the pace of capability improvement is unusually rapid. In a standard leasing market (automobiles, office equipment), the lessor can set multi-year terms that lock in pricing across upgrade cycles. In AI, capability jumps arrive faster than contract cycles, and customers renegotiate accordingly. Third, the relevant "durable good" is not the subscription itself but the knowledge-processing capability it provides. A user who has spent a year working with an AI assistant has developed prompting skills, workflow integrations, and tacit understanding of the system's strengths and limitations. This human capital is partially transferable to competing systems, creating the same downward price pressure that the conjecture predicts.

6 The Coase Theorem and Knowledge Reallocation

The Coase theorem [Coase, 1960] provides a third lens. In its efficiency version, the theorem states that when transaction costs are sufficiently low, bargaining will lead to a Pareto-efficient outcome regardless of the initial allocation of property rights. The initial assignment of rights affects the distribution of surplus but not allocative efficiency.

Applied to the knowledge economy, the theorem yields a striking prediction. The current allocation of knowledge assets, including institutional expertise, proprietary workflows, organizational memory, and client relationships, reflects decades of accumulated transaction costs. Firms that invested in building internal knowledge hierarchies did so because the cost of transferring that knowledge across firm boundaries was prohibitive. This history of costly knowledge transfer created path dependencies: the current distribution of expertise across firms is a product of historical transaction costs, not necessarily of efficient allocation.

As AI drives the transaction costs of knowledge transfer toward zero, Coase's theorem predicts that the initial allocation of knowledge assets becomes less determinative of efficient outcomes. Knowledge can be reorganized across firm boundaries at low cost. A solo practitioner armed with AI can access and deploy expertise that previously required a large institutional infrastructure. A startup can replicate the analytical capabilities of an established consulting firm's research division. The question of who "owns" institutional knowledge becomes less economically meaningful when the cost of replicating or transferring that knowledge approaches zero.

Consider a concrete example. A mid-size law firm's competitive advantage historically rested on the accumulated expertise of its partners in, say, pharmaceutical patent litigation. This expertise was expensive to develop, difficult for clients to evaluate, and nearly impossible to transfer to a competing firm or an independent practitioner. The high transaction costs of knowledge transfer created a durable rent. Under AI, the calculus shifts. An AI system trained on the full corpus of pharmaceutical patent case law can provide a solo practitioner with analytical capabilities that previously required a team of associates and a decade of institutional learning. The transaction costs of knowledge transfer collapse, and the Coase theorem predicts that the efficient allocation of this work will shift from large firms to a more distributed market structure, regardless of where the expertise was historically housed.

This does not mean that all knowledge becomes freely available. The Coase theorem requires low transaction costs, and several categories of knowledge retain high transfer costs even under AI: knowledge that is generated by unique observational vantage points (proprietary data), knowledge that is embedded in trust relationships (client confidence), and knowledge that is protected by legal barriers (regulatory licensing, trade secrets). These residual frictions define the boundaries of the Coasian prediction and identify the durable sources of competitive advantage.

The theorem also carries a distributional warning. Even when bargaining leads to efficient outcomes, the distribution of surplus depends on the initial allocation of rights. If AI enables efficient knowledge reallocation but the gains accrue primarily to those who control AI infrastructure, distribution, and regulatory access, the efficiency gains coexist with increased inequality, a pattern consistent with historical experience with general-purpose technologies [Acemoglu, 2002].

7 Empirical Evidence

The theoretical framework generates testable predictions. In this section, we review early empirical evidence on AI's labor market impacts, drawing primarily on Massenkoff and McCrory [2026], and assess the evidence for and against the Coasian predictions.

7.1 Observed Exposure and Employment

Massenkoff and McCrory [2026] introduce a measure of "observed exposure" that combines theoretical AI capabilities (which tasks AI could perform) with actual usage data (which tasks AI is performing in practice). This distinction is important: theoretical exposure, as measured by Eloundou et al. [2023], captures potential disruption, but actual adoption depends on complementary factors including software integration, regulatory constraints, and verification requirements.

The gap between theoretical and observed exposure is large. For computer and mathematical occupations, theoretical exposure reaches 94% of tasks, but observed coverage is only 33%. This gap is consistent with the GPT framework's prediction of slow complementary adjustment: the technology is available, but the organizational, regulatory, and institutional adaptations required to deploy it fully lag behind.

The key empirical finding is that occupations with higher observed exposure show weaker projected employment growth through 2034 in Bureau of Labor Statistics projections. For every 10 percentage point increase in observed coverage, BLS employment growth projections decline by 0.6 percentage points. This relationship appears only when using observed (actual) exposure, not theoretical exposure alone, suggesting that realized adoption, not merely potential, drives labor market adjustment. Table 2 summarizes the principal findings.

Table 2: Summary of early empirical evidence on AI labor market exposure

Finding	Source
Theoretical exposure: 94% (computer/math); observed: 33%	Massenkoff and McCrory (2026)
-0.6 pp employment growth per 10 pp observed coverage	Massenkoff and McCrory (2026)
No significant unemployment increase for top-quartile exposed workers	Massenkoff and McCrory (2026)
$\approx$14% decline in job-finding rates, ages 22–25, exposed occupations (post-2024)	Massenkoff and McCrory (2026)
$\approx$30% of workers have zero AI exposure	Massenkoff and McCrory (2026)
Exposed workers: older, female, more educated, higher-paid	Massenkoff and McCrory (2026)

7.2 The Displacement Puzzle

A notable feature of the early evidence is the absence of mass displacement. Massenkoff and McCrory [2026] find no systematic increase in unemployment for highly exposed workers since late 2022. Their difference-in-differences framework comparing the top quartile of observed exposure to workers with zero exposure finds no statistically significant divergence in unemployment rates.

This finding is consistent with two interpretations. Under the Coasian framework, the prediction is reorganization rather than displacement: tasks migrate across firm boundaries, workers shift from information processing to judgment-intensive functions, and the transition unfolds gradually as complementary institutional adjustments catch up. Under this interpretation, the absence of mass unemployment is expected in the early phase of a horizontal technology diffusion.

An alternative interpretation, advanced by practitioners and commentators, holds that AI-driven productivity gains increase demand for labor rather than displacing it. The argument, consistent with the Jevons paradox, is that when AI makes workers more productive, firms expand output, enter new markets, and ultimately hire more, not fewer, workers. Proponents point to rapid growth in AI-adjacent job categories and to historical patterns in which productivity-enhancing technologies created more jobs than they destroyed [Autor, 2015].

The task-based framework of Acemoglu and Restrepo [2019] provides a useful formalization. In their model, automation displaces labor from existing tasks, but new technologies also create new tasks in which labor has a comparative advantage. The net employment effect depends on the balance between displacement (automation of existing tasks) and reinstatement (creation of new tasks). The Coasian framework developed here complements this task-level analysis by adding an organizational dimension: AI does not merely automate tasks within existing firm structures but reorganizes which tasks are performed inside firms and which are performed through market transactions. The displacement and reinstatement of tasks occurs simultaneously with the dissolution and reconstitution of firm boundaries.

These interpretations are not mutually exclusive. The Coasian framework predicts both reorganization (shifts in where work is performed and how firms are structured) and a net employment effect that depends on the relative magnitude of displacement and demand expansion. The early empirical evidence is consistent with a transition period in which reorganization dominates displacement, but the long-run equilibrium remains indeterminate.

7.3 Heterogeneous Effects

The distributional pattern in the early evidence aligns with the Coasian prediction. The workers most exposed to AI tend to be older, female, more educated, and higher-paid [Massenkoff and McCrory, 2026]. This is precisely the profile of knowledge workers whose jobs were internalized within firms because their skills were difficult to specify and verify externally: paralegals, financial analysts, medical record specialists, market research analysts.

The one area where early evidence suggests genuine disruption is among young workers entering exposed occupations. Massenkoff and McCrory [2026] find a roughly 14% decline in job-finding rates for workers aged 22–25 in highly exposed occupations after 2024, with no comparable effect for workers over 25. This is consistent with the Coasian prediction that entry-level information-processing roles, the traditional on-ramp into knowledge-economy careers, are the first to be displaced by AI-mediated market transactions.

8 The Residual Firm

If AI dissolves the information asymmetry rents that justified large knowledge-economy firms, what organizational form remains? We characterize the residual firm: the entity that persists after information-processing rents have been competed away.

8.1 Durable Sources of Value

Four categories of competitive advantage survive the Coasian dissolution of information rents.

Distribution and trust. Owning the customer relationship and the trust associated with it is not an information asymmetry rent and is not dissolved by AI. Platform incumbents with large user bases retain structural advantages even as the software and knowledge processing beneath them are commoditized. Trust is built through repeated interaction, reputation, and accountability, none of which AI can substitute for at present.

Proprietary data. Data generated by unique institutional processes, customer relationships, or observational vantage points remains a genuine moat. AI is trained on data; it cannot synthesize proprietary data from public sources. The firm as a collector and curator of unique data retains Coasian justification.

Regulatory position. The frictions AI cannot dissolve are those encoded in law and regulation. Jurisdictional licensing, compliance regimes, and incumbent-friendly regulation become more valuable as other barriers fall. This creates an ironic dynamic in which regulatory capture, historically a source of inefficiency, becomes one of the few durable competitive advantages.

Judgment and problem framing. The non-informational component of high-end professional work, the judgment that integrates values, aesthetics, relationship intuition, and contextual wisdom, is not reducible to information processing. AI answers questions; humans must still generate the right questions. The firm as a locus of judgment, the entity that decides what to optimize rather than how to optimize, retains its Coasian rationale. This is Polanyi's tacit knowledge applied to organizational direction: the firm knows more than it can codify.

8.2 The Structure of the AI Firm

The organizational form that emerges from this analysis is a small nucleus of high-judgment principals with AI-augmented operational reach, dynamically assembling external capabilities as needed rather than internalizing them permanently. Economic rents migrate from information processing to direction-setting and trust. The leverage that matters is not the number of knowledge workers a firm employs but the quality of the judgment at its center and the strength of the relationships at its edges.

Table 3 summarizes which traditional firm rationales survive the AI transition.

Table 3: Coasian rationales for the firm under AI

Traditional Rationale	AI Impact	Status
Reduce search and contracting costs	AI near-eliminates	Weakened
Manage bounded rationality	AI expands individual capacity	Weakened
Capture tacit knowledge	Partially codifiable via AI	Partially weakened
Bear legal and moral accountability	AI cannot be held liable	Survives
Coordinate tacit social capital	AI cannot replicate relationships	Survives
Frame novel problems	AI cannot identify unknown unknowns	Survives

9 Discussion

The Coasian framework developed in this paper unifies several apparently distinct phenomena: the restructuring of professional services, the contraction of knowledge-economy firms, the collapse of software terminal value, and the heterogeneous labor market effects of AI. The common mechanism is the dissolution of information asymmetry rents, and the common analytical structure is Coase's insight that organizational form is endogenous to transaction costs.

Several limitations of the analysis deserve emphasis. First, the framework is deliberately focused on the knowledge economy. For occupations involving physical manipulation, sensory judgment, or in-person service delivery, AI's impact on transaction costs is qualitatively different and less immediate. The Coasian analysis applies most directly to the roughly 70% of workers with nonzero AI exposure identified by Massenkoff and McCrory [2026]; the approximately 30% with zero exposure (cooks, plumbers, electricians, bartenders) fall outside its scope.

Second, the framework treats the pace of AI capability improvement as exogenous. In practice, the rate of improvement is itself influenced by economic incentives, regulatory choices, and the organizational dynamics this paper analyzes. A full general-equilibrium treatment, in which AI capability, firm structure, and labor markets co-evolve, remains an open theoretical challenge.

Third, the distributional implications are ambiguous and carry historical precedent. Ricardo's distinction between labor and capital shares, and the Marxian analysis of how machinery shifts bargaining power from workers to capital owners, find a direct parallel in the AI transition. The "Engels pause" of roughly 1790–1840 is instructive: during the first half-century of British industrialization, GDP grew substantially but real wages stagnated, with the productivity surplus accruing to factory owners rather than workers [Allen, 2009]. The pause ended only when institutional countervailing forces, including labor organization, factory legislation, and expanded suffrage, redistributed bargaining power. AI presents an analogous risk. If the directing function migrates from human managers to AI systems owned by a narrow class of infrastructure providers, the Ricardian division of output shifts toward capital, and a new Engels pause becomes plausible. The distributional question of AI is not determined by the technology itself but by the institutional and political economy that surrounds it [Acemoglu, 2002].

The historical record offers a cautionary precedent on forecasting. Blinder and Krueger [2013] identified approximately 25% of U.S. jobs as vulnerable to offshoring; a decade later, these jobs maintained healthy employment growth. Government occupational projections, while directionally accurate, add minimal predictive value beyond linear trend extrapolation [Massenkoff and McCrory, 2026]. The Coasian framework developed here is a structural theory of direction, not a quantitative forecast of magnitude or timing.

The core claim of this paper is that Coase's intellectual framework, developed in 1937 to explain why firms exist at all, provides the most natural lens for understanding AI's economic impact. AI is a transaction cost annihilator. It dissolves the informational frictions that justified the scale of knowledge-economy firms, the intermediation rents of professional services, the switching costs of enterprise software, and the pricing power of information monopolists. What remains after these rents are competed away, the residual firm, is defined not by its capacity to process information but by its capacity to generate trust, bear accountability, frame problems, and exercise judgment. The irony Coase would appreciate is that the institution he explained by reference to friction is being reshaped by a technology that eliminates friction. What persists is whatever cannot be specified, verified, or monitored: the irreducibly human core of economic organization.

References

Daron Acemoglu. Technical change, inequality, and the labor market. *Journal of Economic Literature*, 40(1):7–72, 2002.

Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.

Armen A. Alchian and Harold Demsetz. Production, information costs, and economic organization. *American Economic Review*, 62(5):777–795, 1972.

Robert C. Allen. Engels' pause: Technical change, capital accumulation, and inequality in the British industrial revolution. *Explorations in Economic History*, 46(4):418–435, 2009.

David H. Autor. Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30, 2015.

Alan S. Blinder and Alan B. Krueger. Alternative measures of offshorability: A survey approach. *Journal of Labor Economics*, 31(S1):S97–S128, 2013.

Timothy F. Bresnahan and Manuel Trajtenberg. General purpose technologies: "Engines of growth"? *Journal of Econometrics*, 65(1):83–108, 1995.

Erik Brynjolfsson and Lorin M. Hitt. Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4):23–48, 2000.

Alfred D. Chandler. *The Visible Hand: The Managerial Revolution in American Business*. Harvard University Press, Cambridge, MA, 1977.

Ronald H. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.

Ronald H. Coase. The problem of social cost. *Journal of Law and Economics*, 3:1–44, 1960.

Ronald H. Coase. Durability and monopoly. *Journal of Law and Economics*, 15(1):143–149, 1972.

Paul A. David. The dynamo and the computer: An historical perspective on the modern productivity paradox. *American Economic Review*, 80(2):355–361, 1990.

Robert J. DeFillippi and Michael B. Arthur. Paradox in project-based enterprise: The case of film making. *California Management Review*, 40(2):125–139, 1998.

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130*, 2023.

Irving John Good. Speculations concerning the first ultraintelligent machine. In Franz L. Alt and Morris Rubinoff, editors, *Advances in Computers*, volume 6, pages 31–88. Academic Press, 1965.

Friedrich A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4): 519–530, 1945.

Michael C. Jensen and William H. Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360, 1976.

Alec Litowitz, Nick Polson, and Vadim Sokolov. The economics of knowledge in the age of LLMs. Technical report, SSRN, 2026a. URL `https://ssrn.com/abstract=6302283`. SSRN 6302283.

Alec Litowitz, Nick Polson, and Vadim Sokolov. Photons = tokens: The physics of AI and the economics of knowledge. Technical report, SSRN, 2026b. URL `https://ssrn.com/abstract=6265418`. SSRN 6265418.

Maxim Massenkoff and Peter McCrory. Labor market impacts of AI: A new measure and early evidence. Technical report, Anthropic, 2026. URL `https://www.anthropic.com/research/labor-market-impacts`.

Michael Polanyi. *The Tacit Dimension*. Doubleday, Garden City, NY, 1966.

Herbert A. Simon. A formal theory of the employment relationship. *Econometrica*, 19(3): 293–305, 1951.

Oliver E. Williamson. *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press, New York, 1975.

Oliver E. Williamson. *The Economic Institutions of Capitalism*. Free Press, New York, 1985.