Transformative AI and Firms

Aaron Chatterji
Duke University & OpenAI

Daniel Rock
University of Pennsylvania

Eduard Talamàs
IESE Business School

October 1, 2025

Abstract

We discuss key economic questions surrounding firms in a world where "transformative AI" (TAI)—described as a "datacenter full of geniuses" in Amodei (2024)—exists in every domain. We explore how firms may be able to adapt, compete, and evolve in this scenario. We propose a series of theoretical research questions and associated measurement challenges. Firms play a critical role in the adoption and diffusion of AI tools. We discuss open questions about potential firm structures, investment decisions, and competitive implications that could result from the advent of transformative AI. Firms have proven to be powerful aggregators of information and essential to producing and organizing knowledge, ultimately driving economic growth. TAI may fundamentally change the role of companies and how they are organized, shaping the impact of TAI on the economy.

Keywords: Transformative Artificial Intelligence, Intangible Capital, Competition, Productivity

1 Introduction

Considering how firms might be structured or behave differently in the presence of extraordinarily capable, abundant digital intelligence is a challenging exercise. Leaders of the top artificial intelligence companies and laboratories advise that some version of artificial general intelligence (AGI) is imminent (Amodei, 2024; Altman, 2025). In the spirit of this volume, we will bypass the precise definition of AGI in this paper to discuss "transformative" AI (TAI): a set of artificial intelligence capabilities that are assumed to be 1) easily proliferated and 2) "intelligent" enough to surpass human intellectual capacities in most domains. In short, we take Dario Amodei's question seriously: what if anyone could access a "datacenter full of geniuses"? If this comes to pass, it will have profound economic implications as machine capabilities expand rapidly, both enabling impressive value creation and imposing enormous risks to existing social and economic structures. We set aside questions about the likelihood of this scenario. Assuming TAI *is* imminent, we discuss some of what could be important to understand about how firms could operate in this context.

Companies are prosaic objects of study if intelligent software is abundant, but the firm is the primary means of coordinating and sharing the creation of economic value in the modern economy. If human beings find advantages in corporations, perhaps humans collaborating with transformative AI agents will as well. We have an overarching question: what would remain constant and what would change for firms in a TAI-enabled world? This article discusses some relevant frameworks and research questions to pursue as part of a TAI economics agenda, focusing on the firm as an important operating unit in a TAI economy. As the technological change under consideration is by definition transformative, the items discussed are non-exhaustive. We mean to offer a few starting points to stimulate future research.

We first discuss a set of questions related to firm structure and production in section 2. We then consider the role of firms in the organization and production of knowledge in section 3. Section 4 concludes by posing a series of research questions related to competition and demand-side TAI changes.

2 Supply-side Considerations: Productivity, Labor, and Capital

In "Machines of Loving Grace", Dario Amodei suggests that we should be thinking about the marginal returns to intelligence once these genius software capabilities proliferate (Amodei, 2024). Accordingly there are a number of complements to transformative artificially intelligent agents Amodei recommends considering, including speed of the outside world, need for data, intrinsic complexity, constraints from humans, and physical laws. These are primarily environmental factors when it comes to firm behavior, though firms might configure better means of processing data (for example). Let's instead consider intelligence as an input in a standard production function linking output Y to capital K, labor L, total factor productivity A, and a new input I that can either increase total factor productivity or serve as an ordinary production input:

$$Y = A(I) * F(K, L, I) \tag{1}$$

If we have competitive factor markets, the rental rate of capital r, the wage w, and the "rental rate of TAI intelligence" r_I reflect their marginal products. In the case that there are substantial adjustment costs of investment in any of the input factors, the difference between the marginal adjustment costs of competitors and those of a focal firm will be appropriable as quasi-rents in the short-run for the more efficient company. That is, they will make some money. TAI under these circumstances might be an important input for companies to create and market enormously valuable products, but the underlying economic toolkit to understand these firms would not need to change. The challenges would primarily be empirical, and the TAI entities available in the datacenter could help economists to measure their own impact.

TAI is distinctive in part because it represents the digitization of genius-level labor—a form of work that may have long been a major bottleneck in the economy (Benzell and Brynjolfsson, 2019). Scarce and expensive talent on the margin is converted to a high fixed cost, very low (or zero) marginal cost input. Given the commodification of intelligence, incumbent firms would have strong incentives to restructure their workflows, while new entrants could design entirely different processes from the ground up. We consider implications for firm labor demand, capital accumulation, and total factor productivity below.

2.1 Standard Inputs to the Production Function

Even current vintages of AI are changing the labor market in pervasive ways, with the potential for even more widespread changes (Felten et al., 2023; Eloundou et al., 2024; Bick et al., 2024; Handa et al., 2025). Recent research indicates knowledge workers tend to be more exposed to large language models than people doing physical work. TAI would likely be no different on the extensive margin, with radically increased potential on the intensive margins of exposure. Whether measured as "exposure" or actual usage in specific work contexts, new technologies can be either beneficial or harmful to the incumbent workforce. However, their introduction invariably introduces risk by making previously stable outcomes more volatile.

In a world of abundant digitized intelligence, we can consider the limiting case in which the marginal return to intelligence I—and thus its competitive price—approaches zero. This could make it difficult for some types of labor to earn wages. Knowledge work of today's varieties could be radically repriced, with machines substituting for people in many circumstances. Following the task-based model of automation (e.g., Acemoglu and Restrepo, 2018), TAI might substitute for some tasks of knowledge workers and allow the employer firm to substitute them with capital. The risk to human capital investment increases rapidly as some workers are made more productive by TAI (e.g., for workers with expertise that cannot be accessed by the AI systems) and other types of work become obsolete. Expertise may be an important determinant of labor demand (Autor and Thompson, 2025). Firm-specific expertise that enables workers to add context to TAI workstreams might be particularly valuable as general TAI capital expands into more tasks. This expertise is a form of capital for employer firms, albeit with imperfect alienability from the employees. In many models of production, labor is assumed to be a flexible input while capital has fixed costs of adjustment (e.g., Lucas Jr, 1967). In the presence of transformative AI, some intangible human capital becomes marketable and flexible machine capital. In a TAI scenario we may have to revisit the assumptions that fixed costs of labor adjustment are negligible and that knowledge capital adjustment (of some varieties) is expensive.

Concurrently there could be an incentive with TAI to use this new-found genius capital to endogenously create new tasks for labor. In traditional firms this is one of management's responsibilities. New tasks "reinstate" labor, often in complementary formats to other factors of production (Acemoglu and Restrepo, 2018). The extent and speed to which TAI leads to new work for humans—as well as its impact on capital accumulation, profits, and business models—will be important directions for empirical research. In these task-driven models (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018), new technologies can change total factor productivity, alter factor-specific productivities, enable substitution between capital and labor varieties, and lead to creation of new tasks.

Whether worker tasks are automated or augmented is often cited as the core tension for labor demand studies (Brynjolfsson, 2022). The difference is less useful for specific firms and their employees than it is for discussing employment overall, as the shape of labor demand governs wages, employment, and other economic outcomes. Automation of some tasks with TAI may lead to an expansion of demand for residual tasks in the same job, or it might lead to layoffs. Augmenting workers to be more productive can raise wages and expand employment with elastic demand, or it might enable one worker to replace many of their colleagues if demand is inelastic.

One core difference with TAI relative to other digital technologies is the potential to impact executive functions within firms. Decision-making power is traditionally centralized in modern corporations. We detail in section 3 that economic considerations change when AI can generate and act upon knowledge in a decision-making capacity. Importantly, tasks that can be framed with an objective function are more readily delegated to machine systems. The measurable work becomes the machine-exposed work (Catalini et al., 2025). Since the most measurable tasks might become the most productive tasks, the limiting factors of production might be in areas where AI has the least influence, resembling a version of Baumol's cost disease (Aghion et al., 2017).

Capital in modern firms is configured for deployment by human decision-makers. Similar to human capital and labor, the economic value of firm capital stocks will be repriced in the presence of TAI. If TAI deployment is capital-intensive, we would expect the real cost of capital to increase rapidly (Chow et al., 2024). TAI would potentially drive the creation of new capital tasks and increased capital productivity. Each of the factors in the task model apply equally well to capital as they do to labor, though perhaps capital effects are less often discussed. Obsoleting and replenishing capital stocks across a wide swathe of companies will be expensive and subject to frictions. One of the critical open questions is which barriers to capital investment in the presence of TAI will slow potential deployment of the technology. Some of these transitional frictions may be principally related to intangible factors rather than tangible capital. We discuss those next.

2.2 Innovation and Intangible Capital with Transformative AI

With transformative AI, the problem where ideas are increasingly expensive to find (Bloom et al., 2020) may no longer be a concern. Firms could be able to easily create new TFP-increasing ideas, though a core challenge with commodified genius is realizing any profits from those ideas. Profit, to the extent it exists, might accrue to other factors. TFP can grow as the familiar A becomes an endogenous function of intelligence (or $A(I)$), as in (Romer, 1990). This leads to another open question for researchers: what are the new scarce factors of production and how long does it take for artificial intelligence to uncover the high value intellectual property of competitors? It may also be the case that the proliferation of ideas creates a high-dimensional competitive idea space that is hard to "fill" with businesses because of scarce production inputs.[1] This could lead to interesting research questions surrounding contracting and intellectual property. Historically IP-related risks have limited the extent to which firms can contract with each other. If, on the other hand, IP is impossible for TAI to hide from other AI-enabled firms, it might change the benefits and costs of outsourcing work since any IP is transparent (or discoverable for a price) to all trading partners. TAI could significantly change the boundary of the firm as well.

[1] This is one scenario where TAI could increase human wages and employment.

An alternative to the task model is a systems-based approach (Bresnahan, 2019; Agrawal et al., 2024). Firms are systems with many component modules, each of which might be affected differently by technology. Corporate decision-makers are responsible for altering the sequence and structure of internal processes. TAI would radically change incentives in making these decisions, open up more potential paths for productive activity, and expose incumbents to new competitors with more efficient AI-enabled organizational designs. Some technology leaders, including OpenAI's Sam Altman, have speculated that a single person unicorn company might be possible in the near future because of AI (Yahoo Finance, 2025).

We might consider then extending the standard production function approach to account for the internal structure of firms. Theoretically, each firm in a systems-based model could be consistent not only of a mapping from inputs to output, but also a set of paths describing the input-to-output map. In other words, firms can be represented as a directed acyclic graph (DAG) that takes a set of input factors and converts them to output. TAI, as with the introduction of any new technology, would create incentives to alter the productive pathways for inputs inside of the firm. These kinds of pathway adjustments are costly, but potentially necessary when facing a competitor with a more efficient DAG. A solo founder with an army of AI employees would constitute a very different organizational structure than an incumbent with thousands of human employees, primarily in how knowledge is created and monetized. Of course, any structure that a solo founder can create is likely to face competition with zero marginal cost TAI.

In terms of a modeling approach, we might consider firms not as single production function nodes as in equation 1, but rather as a network of interdependent pathways of production nodes leading to output. Firms not only optimize within nodes, but then across them, picking the optimal path given prices and a target level of output. Relatively "deeper" firms have more stacked production processes to map from marketed inputs to outputs, but this depth makes adjustment more costly. Firms may attempt to incorporate TAI into existing processes as well. That can increase productivity, but more likely gains are to be found at the end of production processes where there are not as many downstream bottlenecks. Over time, TAI will require the construction of new paths (as a form of organizational capital). Constructing these paths will necessitate the diversion of human capital away from short-term productive activity into organizational capital creation. The market for talent that can create these new production paths might crowd out investment in short-term productive activity with those same workers. Expensive reconfiguration simultaneously offers an entry opportunity for new firms without the same incumbent structures to adjust.

Another key potential difference between TAI and other technological changes is the potential velocity for these competitive dynamics to play out. Typically new types of intangible capital and systems take many years to build and accumulate (Brynjolfsson et al., 2021). Fast-moving digital capital accumulators shorten the clockspeed if AI agents can build their own intangible knowledge capital for coordination. We discuss TAI and research relating to firm knowledge creation in the next section.

3 TAI and Knowledge Work in Firms

Within the firm, TAI technologies could fundamentally alter the way work is structured and the opportunities to solve problems. Transformative AI represents a fundamentally new kind of automation technology: one that enables machines to perform non-codifiable knowledge work (Section 3.1). Below we outline a research agenda on the transformative implications of this development for knowledge firms, focusing on how AI may reorganize knowledge work (section 3.2) and reshape the role of organizations in the production and codification of knowledge (section 3.3).

3.1 AI is Redefining the Boundaries of Automation

Traditional automation has been limited to tasks governed by clearly defined, codifiable rules—such as arithmetic operations or assembly line routines (Autor et al., 2003; Autor, 2014). In other words, if we couldn't articulate how to do something in precise steps, we couldn't automate it.

This codifiability constraint sharply limited the scope of automation, excluding a wide range of tasks that rely on tacit knowledge—skills that are intuitive, experiential, and hard to put into words. A senior lawyer, for instance, may develop an instinct for complex cases over years of practice, yet struggle to explain or formalize that intuition. As Polanyi (1966) famously put it, "we know more than we can tell."

AI marks a significant shift in the nature of automation because it breaks this traditional codifiability constraint (e.g., Brynjolfsson and McAfee, 2014; Brynjolfsson and Mitchell, 2017; Brynjolfsson et al., 2018; Autor, 2024; Ide and Talamàs, 2025). Machine learning systems learn from data rather than being programmed step by step, uncovering patterns and rules that neither humans nor machines can easily articulate. In this sense, machines are becoming more human-like: they, too, now "know more than they can tell."

This fundamental breakthrough is allowing machines to perform work that until very recently was widely believed to be the exclusive domain of humans—including coding, planning, driving, and conducting research—opening up seemingly limitless possibilities for automation. As these capabilities expand, technology companies are no longer merely delivering software tools; they are developing and deploying powerful new digital agents that are bound to transform the knowledge economy.

3.2 Transformative Impact via Reorganization of Knowledge Work

As emphasized by Ide and Talamàs (2025), AI's ability to perform non-codifiable knowledge work gets at the heart of a central bottleneck identified by the literature on knowledge hierarchies, beginning with Garicano (2000).[2] The central ideas in this literature are that (i) the practical use of tacit knowledge is constrained by the limited time of the experts who possess it, and (ii) organizations play a major role in alleviating this constraint.

From this perspective, whether on a factory floor, in a law firm, or in a hospital, organizational design rests on a simple principle: shield those with valuable, specialized knowledge from tasks that do not require their unique expertise. This is achieved by structuring communication flows so that only the most complex or exceptional problems reach the experts. This principle—known as management by exception—is captured by the former head of General Motors, Alfred Sloan (1924), who remarked: "We do not do much routine work with details. They never get up to us. I work fairly hard, but on exceptions."[3]

The introduction of AI agents is set to transform the knowledge economy by enabling machines to perform work that relies on tacit, non-codifiable knowledge (Ide and Talamàs, 2025). This undermines the bottlenecks that have long shaped traditional knowledge hierarchies, prompting a fundamental reorganization of knowledge work.

Ide and Talamàs (2025) introduces AI agents into a Garicano (2000)-style model of the knowledge

[2] For a comprehensive overview, see Garicano and Rossi-Hansberg (2015).

[3] See Garicano and Hubbard (2012); Caliendo and Rossi-Hansberg (2012); Caliendo et al. (2015, 2020) for recent empirical evidence.

economy, where humans form hierarchical organizations to solve problems. Their analysis builds on the baseline framework developed by Garicano and Rossi-Hansberg (2004), Antràs et al. (2006), and Fuchs et al. (2015)—in which labor is the sole factor of production.

In this pre-AI framework, humans are endowed with one unit of time and differ in their knowledge. Individuals pursue production opportunities, encountering problems of varying difficulty in the process. They succeed only when their knowledge exceeds the difficulty of the problems they encounter. To make the best possible use of the available time and knowledge, organizations construct knowledge hierarchies: when a worker is unable to solve a problem, she can escalate it to a more knowledgeable individual higher up in the hierarchy.

Ide and Talamàs (2025) models AI as a technology that converts computing power into agents that can operate either autonomously—pursuing projects independently—or non-autonomously, assisting with problem-solving but unable to complete projects on their own. This distinction shapes how AI reorganizes the knowledge economy: autonomous AI boosts aggregate output more but widens labor income inequality by favoring the most knowledgeable, while non-autonomous AI benefits less knowledgeable workers but yields lower overall output.

3.3 Transformation via Changes in Knowledge Production

The emergence of AI agents capable of performing non-codifiable work autonomously and at scale raises foundational questions not only about how firms can leverage existing knowledge, but also about the role of organizations as engines of knowledge creation and codification. We now outline several key questions that arise in this latter domain.

3.3.1 A New Type of Knowledge: Machine-Tacit Knowledge

Traditional *human-tacit knowledge*—the things we know but can't tell—now has a new counterpart: *machine-tacit knowledge*, or the things machines know but can't tell (Ide and Talamàs, 2025). These two forms of tacit knowledge are fundamentally different: Human-tacit knowledge is difficult to articulate and transfer—both to other humans and machines—and thus remains dispersed across individuals. In contrast, machine-tacit knowledge is embedded in model weights, making it easily transferable between machines. This transfer requires neither explanation nor interpretation; it occurs simply through the copying of weights.

Thus, AI raises foundational questions about the future relevance of knowledge codification and the evolving role of organizations in that process. To explore these questions, we begin by revisiting the traditional codifying role of organizations.

3.3.2 The Traditional Knowledge Codifying Role of Organizations

The importance of organizations in the process of knowledge codification has long been recognized (e.g., Simon, 1947; Argyris and Schön, 1978; Nonaka, 1991). In his classic essay, Nonaka (1991) provides a vivid illustration: In 1985, engineers at Matsushita Electric Company set out to create a home bread-making machine—but the dough kept baking unevenly. Exhaustive X-ray comparisons of machine-kneaded and artisan-kneaded loaves produced no actionable insight; the team simply could not automate the baker's touch.

Software developer Ikuko Tanaka proposed a radical remedy: become an apprentice. She spent months at the Osaka International Hotel, watching the head baker's distinctive stretch-and-twist motion and practicing it herself. Back at Matsushita, Tanaka and the engineers iterated for a year, ultimately embedding that motion in the machine via the "twist-dough" method—allowing it to deliver loaves indistinguishable from the baker's.

The impact of Tanaka's work didn't stop with the twist-dough method. By articulating the baker's tacit knowledge into an explicit design, she triggered what Nonaka calls "the spiral of knowledge"—a dynamic process in which explicit knowledge circulates, gets internalized by others, and then reinterpreted to spark further innovation. The codified kneading technique became a foundation upon which new generations of engineers could build, improve, and combine with other forms of knowledge. In this way, one individual's embodied insight was transformed into an organizational asset—reusable, teachable, and expandable.

3.3.3 Knowledge Creation, Codification and Interpretability in the Age of AI

The fact that machine-tacit knowledge can be costlessly transferred between machines—simply by copying model weights—may diminish the value of codification. Yet it may just as well increase it if making such knowledge intelligible and actionable for humans becomes increasingly important. This raises a pressing question: can we design organizations of AI agents that actively engage in codifying their own tacit knowledge? And what role—if any—can humans play in shaping or interpreting that process?

On this point, Nonaka (1991) offers further valuable guidance. In human organizations, knowledge creation thrives when diverse perspectives are brought into dialogue, when a shared language allows people to refine and reframe ideas, and when metaphors and analogies translate experience into shared understanding. These mechanisms help individuals externalize tacit knowledge and transform it into organizational assets. Might similar processes be needed to codify machine-tacit knowledge? And if so, might humans still be able to contribute to this process—at least on the diversity front?

These questions connect to the growing literature on machine interpretability, which aims to make the internal representations and decision-making processes of AI systems accessible to human reasoning. Techniques such as feature attribution and model distillation serve to render opaque model behavior into forms humans can understand and act upon (Lipton, 2018; Doshi-Velez and Kim, 2017; Gilpin et al., 2018). In this light, interpretability can serve as a form of codification, one that may allow hybrid teams of humans and machines to collaborate in the all-important pursuit of innovation and knowledge creation.

Even as AI systems grow vastly more capable, interpretability may remain essential—not because intelligence is scarce, but because codifiable knowledge is. As Amodei (2024) puts it: "In the AI age, we should be talking about the marginal returns to intelligence, and trying to figure out what the other factors are that are complementary to intelligence and that become limiting factors when intelligence is very high." Codifiable knowledge may be one of those critical complementary factors. If so, interpretability methods may become the scaffolding for collective sense-making in organizations.

3.3.4 Implications for Theories of Firm, Strategy and Entrepreneurship

The dawn of TAI would require a reconsideration of many of the leading theories of the firm. These influential theories have not only spawned a large literature in economics but has been foundational in other areas of study, notably strategic management and entrepreneurship.

Most theories of the firm begin with the question of why firms exist and what factors shape their boundaries. The property rights theory of the firm (Grossman and Hart, 1986; Hart and Moore, 1990) argues for the central role in asset ownership in dictating incentives and determining firm boundaries. Another seminal stream of work on transaction costs economics (Williamson, 1975,

1996) places more emphasis on the characteristics of transactions themselves and the role of firms in mitigating contractual hazards in markets. TAI could usher in an era of AI infrastructure becoming the most valuable of all assets and thinking machines engaging in the vast majority of transactions in the economy. These developments would alter the calculus around asset ownership and arguably make the risk of opportunism less relevant, provoking a reconsideration of these two important theories of the firm.

Other important theories emphasize the role of organizational governance to enable monitoring and reduce free-riding in team production (Alchian and Demsetz, 1972) and cast firms as the most efficient way to organize uniqye and valuable knowledge (Kogut and Zander, 1992).These theories would also have to be revised to consider how widely available access to an essentially unlimited stock of knowledge would influence both the relative importance of monitoring human effort and the nature and amount of knowledge that could be efficiently organized within firm boundaries. Interestingly, as humans and machines increasingly work together in organizations, the insights from more behavioral theories of the firm that consider the importance of relational contracts and internal politics might become even more applicable (Cyert and March, 1963; Baker et al., 2002).

Many of these theories of the firm have influenced the academic literature on business strategy, which focuses on explaining persistent variation in firm performance. As TAI is adopted by firms, it could create significant variation in performance. For example, a large literature on technology adoption in firms finds that the adoption of specific technologies like information technology and complementary management practices can drive performance benefits (e.g., (Bloom et al., 2012)). Most studies of industry technology adoption acknowledge frictions and adjustment costs, leading some firms to adopt technologies before others. This may be because these firms have the appropriate capabilities to incorporate new technology more effectively into their operations or because they lack legacy systems that inhibit adoption. This uneven pattern of adoption and diffusion across firms can reinforce competitive advantage within industries or upend the existing market structure via entrepreneurial entry.

While incumbent firms have sometimes been slow in adopting new innovations, the diffusion of transformative AI is not straightforward to predict. On one hand, new startups without existing workflows might find it easier to adopt TAI while large incumbents struggle to infuse TAI into their legacy workflows. On the other hand, adopting TAI might require complementary assets like access to compute and relationships with large foundation model providers that favor larger organizations.

TAI can also enable the strategies that firms develop to pursue market opportunities. For example, TAI might have significant implications for how firms select the markets that they compete in and how they organize. As AI systems grow in capabilities, managers will be able to delegate some of the strategy-making process to machines. Market and competitor analysis and scenario planning could all be enabled by AI. However, if all firms have access to this technology, a differentiating factor could be in what internal data the firm has and what kind of model they train. Another differentiating factor could be in how firms use TAI to reorganize their business units, shedding costs by using AI for functions like human resources and legal or improving coordination between divisions and lowering communication costs.

The existence of TAI will also impact supply and demand conditions that shape firm performance. For example, if a concentrated set of firms controls TAI upstream, it will raise costs and lower margins for firms downstream who use the technology in their organizations or in their products. If TAI provides buyers greater transparency around pricing, firms could also see margin compression. Organizations stuck in the middle of a value chain book-ended by concentrated upstream compute and consumers with significant bargaining power would likely have to reconsider their business model and firm boundaries.

TAI will also have significant implications for entrepreneurship. New technologies typically create opportunities for new firms and business models. TAI will impact firm creation across the entire value chain. There is projected to be tremendous demand for new energy demand and hardware and software for data centers. Many companies will emerge to optimize these parts of the value chain. There will also be numerous companies who build "on top of" foundation models to provide verticalized offering in finance, health, energy and other key sectors.

TAI may also shift the size distribution of entering firms. One possibility is startups can be more capital-efficient by leveraging TAI. If firms can do more with less, venture capitalists could place more bets from their investment funds, potentially changing financing considerations across industries.

4 Conclusion

We conceptualize Transformative AI (TAI) as a new input to production that digitizes genius labor, introducing both immense potential for productivity gains and challenges in labor market adjustment. With near-zero marginal costs, TAI could drastically reprice knowledge work, altering labor demand, task composition, and firm workflows while raising questions about automation, augmentation, and human capital obsolescence. Economically, TAI reshapes traditional models by affecting the marginal returns to intelligence, the pricing of firm capital, and total factor productivity, while also complicating profit realization due to IP transparency and idea saturation. Firms will likely need to reorganize around these changes, with some possibly emerging as lean, AI-first entities optimized around new digital capabilities. The future of firm organization, labor allocation, and capital deployment under TAI will hinge on how firms reconfigure internal processes, adapt to technological complementarities, and navigate emerging empirical and theoretical challenges.

As AI begins to perform tasks that—until very recently—were assumed to rely on human tacit knowledge, the boundary between codifiable and non-codifiable work is being blurred. This transformation raises fundamental questions about how firms generate, structure, and apply knowledge. Making sense of this shift demands renewed attention to the architecture of knowledge work and to the evolving role of firms as engines of innovation and codification. The research agenda ahead is foundational to understanding—and shaping—the future of knowledge in the age of TAI.

Our discussion of the effects of TAI on firms has focused primarily on supply-side considerations, leaving aside the structure of market demand. Yet in a TAI scenario we are just as likely to have extensive adoption of AI in consumer applications. TAI would not tire of bargain hunting, it could learn preferences to represent or predict human buyer interests, or even have consumer interests on its own. TAI consumers could extend some existing trends. For example, if production shifts toward more digital goods, recommendation systems and discovery of niche content could be significantly improved with AI agents carrying preferences for human consumers. At the same time, it's possible the composition of demand shifts meaningfully toward whatever a TAI populace "wants", whether that demand is for its own sake or in service of completing tasks downstream of human interests.

The series of research questions we have explored pertain to a "minimally transformative" AI scenario—one in which intelligence is available at nearly zero marginal cost, yet still falls short of the capabilities of an omniscient planner. This is an important caveat. If such a planner-like AI were to emerge, it could render firms and organizations obsolete as meaningful units of analysis, potentially solving Hayek's problem of knowledge distribution without the need for markets (Hayek, 1945). However even increasingly powerful intelligences may still benefit from coordination within organizational structures. The task of allocating scarce resources among agents with diverse and unbounded preferences remains a complex computational challenge—even for highly advanced systems (Shalizi, 2012).[4] Should such a transformation occur, a new research agenda would be required to understand economic and organizational dynamics under even more powerful AI. Still, just as the study of firms today provides insights into how they might evolve under TAI, analyzing TAI-era firms may, in turn, offer a lens into the possibilities of a more radically transformed future.

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108(6):1488–1542.

Aghion, P., Jones, B. F., and Jones, C. I. (2017). Artificial intelligence and economic growth. Technical report, National Bureau of Economic Research.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2024). Artificial intelligence adoption and system-wide change. *Journal of Economics & Management Strategy*, 33(2):327–337.

Alchian, A. A. and Demsetz, H. (1972). Production, information costs, and economic organization. *The American economic review*, 62(5):777–795.

Alfred Sloan (1924). The most important thing i ever learned about management. *System: The Magazine of Business*, 194(46):140–141.

Altman, S. (2025). Reflections. Accessed: 2025-07-13.

Amodei, D. (2024). Machines of loving grace: How ai could transform the world for the better. October, 2024. Available at https://darioamodei.com/machines-of-loving-grace (accessed October 14, 2024).

Antràs, P., Garicano, L., and Rossi-Hansberg, E. (2006). Offshoring in a knowledge economy. *The Quarterly Journal of Economics*, 121(1):31–77.

Argyris, C. and Schön, D. A. (1978). *Organizational Learning: A Theory of Action Perspective*. Addison-Wesley, Reading, MA.

Autor, D. (2014). Polanyi's paradox and the shape of employment growth. Technical report, National Bureau of Economic Research.

Autor, D. (2024). Applying ai to rebuild middle class jobs. Technical report, National Bureau of Economic Research.

Autor, D. and Thompson, N. (2025). Expertise. Technical report, National Bureau of Economic Research.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333.

Baker, G., Gibbons, R., and Murphy, K. J. (2002). Relational contracts and the theory of the firm. *The Quarterly Journal of Economics*, 117(1):39–84.

[4]Though perhaps not—future advances may reveal new ways of computing equilibria efficiently (Kroer et al., 2019).

Benzell, S. G. and Brynjolfsson, E. (2019). Digital abundance and scarce genius: Implications for wages, interest rates, and growth. Technical report, National Bureau of Economic Research.

Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative ai. Technical report, National Bureau of Economic Research.

Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2020). Are ideas getting harder to find? *American Economic Review*, 110(4):1104–1144.

Bloom, N., Sadun, R., and Reenen, J. V. (2012). Americans do it better: Us multinationals and the productivity miracle. *American Economic Review*, 102(1):167–201.

Bresnahan, T. (2019). Artificial intelligence technologies and aggregate growth prospects. *Prospects for economic growth in the United States*, pages 132–172.

Brynjolfsson, E. (2022). The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence. *Daedalus*, 151(2):272–287.

Brynjolfsson, E. and McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & company.

Brynjolfsson, E. and Mitchell, T. (2017). What can machine learning do? workforce implications. *Science*, 358(6370):1530–1534.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What can machines learn and what does it mean for occupations and the economy? In *AEA papers and proceedings*, volume 108, pages 43–47. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.

Caliendo, L., Mion, G., Opromolla, L. D., and Rossi-Hansberg, E. (2020). Productivity and organization in portuguese firms. *Journal of Political Economy*, 128(11):4211–4257.

Caliendo, L., Monte, F., and Rossi-Hansberg, E. (2015). The anatomy of french production hierarchies. *Journal of Political Economy*, 123(4):809–852.

Caliendo, L. and Rossi-Hansberg, E. (2012). The impact of trade on organization and productivity. *The Quarterly Journal of Economics*, 127(3):1393–1467.

Catalini, C., Wu, J., and Zhang, K. (2025). What gets measured, ai will automate. *Harvard Business Review*. Accessed: 2025-07-14.

Chow, T., Halperin, B., and Mazlish, J. Z. (2024). Transformative ai, existential risk, and real interest rates. Technical report, Working Paper.

Cyert, R. M. and March, J. G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall, Englewood Cliffs, NJ.

Doshi-Velez, F. and Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). Gpts are gpts: Labor market impact potential of llms. *Science*, 384(6702):1306–1308.

Felten, E., Raj, M., and Seamans, R. (2023). How will language modelers like chatgpt affect occupations and industries? *arXiv preprint arXiv:2303.01157*.

Fuchs, W., Garicano, L., and Rayo, L. (2015). Optimal contracting and the organization of knowledge. *The Review of Economic Studies*, 82(2):632–658.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5):874–904.

Garicano, L. and Hubbard, T. N. (2012). Learning about the nature of production from equilibrium assignment patterns. *Journal of Economic Behavior & Organization*, 84(1):136–153.

Garicano, L. and Rossi-Hansberg, E. (2004). Inequality and the organization of knowledge. *American Economic Review*, 94(2):197–202.

Garicano, L. and Rossi-Hansberg, E. (2015). Knowledge-based hierarchies: Using organizations to understand the economy. *Annual Review of Economics*, 7(1):1–30.

Gilpin, L. H., Bau, D., Yuan, B. Z., Bajwa, A., Specter, M., and Kagal, L. (2018). Explaining explanations:

An overview of interpretability of machine learning. In *2018 IEEE 5th International Conference on data science and advanced analytics (DSAA)*, pages 80–89. IEEE.

Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of political economy*, 94(4):691–719.

Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., et al. (2025). Which economic tasks are performed with ai? evidence from millions of claude conversations. *arXiv preprint arXiv:2503.04761*.

Hart, O. and Moore, J. (1990). Property rights and the nature of the firm. *Journal of political economy*, 98(6):1119–1158.

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

Ide, E. and Talamàs, E. (2025). Artificial intelligence in the knowledge economy. *Journal of Political Economy*.

Kogut, B. and Zander, U. (1992). Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization science*, 3(3):383–397.

Kroer, C., Peysakhovich, A., Sodomka, E., and Stier-Moses, N. E. (2019). Computing large market equilibria using abstractions. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 745–746.

Lipton, Z. C. (2018). The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57.

Lucas Jr, R. E. (1967). Adjustment costs and the theory of supply. *Journal of political economy*, 75(4, Part 1):321–334.

Nonaka, I. (1991). The knowledge-creating company. *Harvard Business Review*, 69(6):96–104.

Polanyi, M. (1966). *The Tacit Dimension*. Doubleday & Co.

Romer, P. M. (1990). Endogenous technological change. *Journal of political Economy*, 98(5, Part 2):S71–S102.

Shalizi, C. (2012). In soviet union, optimization problem solves you. Blog post on *Crooked Timber*. Accessed: 2025-07-13.

Simon, H. A. (1947). *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organizations*. Macmillan, New York. 3rd ed., 1976; reissued with new introduction, Free Press, 1997.

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

Williamson, O. E. (1996). *The mechanisms of governance*. Oxford university press.

Yahoo Finance (2025). Could ai create a one-person unicorn? sam altman thinks so — and silicon ... https://finance.yahoo.com/news/could-ai-create-one-person-120000722.html. Accessed: 2025-07-14.