AI's Use of Knowledge in Society

Erik Brynjolfsson*
Zoë Hitzig†

September 2025

Prepared for NBER Volume on the Economics of Transformative AI.

Abstract

Hayek's famous insight was that central planning – even if economically efficient – is not feasible because the necessary knowledge is inherently dispersed throughout the economy. The advent of transformative AI demands a reappraisal of Hayek's argument and its implications for how decision rights are allocated within firms and society. We develop a property rights framework in which powerful AI can shift the optimal locus of control through two channels: (i) by codifying local knowledge that was previously tacit and inalienable, and (ii) by expanding information processing capacity to aggregate, interpret, and act on data. These forces erode the informational advantage of maintaining on-the-spot decision-makers, making centralized coordination and control more feasible and more efficient – especially where complementarities across assets are important. The framework yields several predictions: larger average firm size, greater industry concentration, and reduced local managerial autonomy. We review early evidence and find that it is largely consistent with these patterns. We also discuss the choices and conditions that can still favor decentralization. The implications of our analysis extend beyond economic considerations: centralization of economic power can lead to centralization of political power and dampen incentives to invest in human capital.

*Stanford Digital Economy Lab, Institute for Human-Centered AI and NBER.
†Harvard Society of Fellows.

Acknowledgements: We are grateful to Seth Benzell, Avi Goldfarb, Ben Golub, Andreas Haupt and Andrew Ng for valuable conversations and comments on this topic and the Stanford Digital Economy Lab for funding. Any errors remain our own.

The 'data' from which the economic calculus starts are never for the whole society 'given' to a single mind... and can never be so given.

— F. A. Hayek, 1945

We (the whole industry, not just OpenAI) are building a brain for the world.

— Sam Altman, 2025

1 Introduction

Hayek's 1945 essay "The Use of Knowledge in Society" opens by criticizing a plan for constructing a "rational economic order" from the top down (Hayek, 1945). He was responding primarily to Oskar Lange and Abba Lerner, whose writings of the 1930s and early 1940s argued that a central planner armed with enough information could efficiently direct an entire economic system (Lange, 1936; Lerner, 1938). That is, economic decision-making could be highly centralized: a planning board would set prices and production targets, guiding what producers make and, through those prices, what consumers choose to buy. Lange and Lerner, in turn, were responding to Ludwig von Mises's earlier critiques of socialism (von Mises, 1920, 1935), which argued that, without private ownership of the means of production, no genuine market prices for capital goods can form, making efficient allocation by a central planner impossible. Together, these exchanges formed what became known as the Socialist Calculation Debate.

Hayek's contribution to the debate was epistemic: it is impossible, he argued, to make all the knowledge required for efficient allocations known to a central entity. As he put it, "The 'data'... are never for the whole society 'given' to a single mind... and can never be so given." There are physical limits on which forms of information can travel, and how quickly. Much of what people know and which is relevant to choosing the best economic outcomes, is local or specialized knowledge, tailored to the instant, or otherwise tacit and difficult to articulate. It is impossible, he argued, to gather, codify and transfer that "knowledge of the particular circumstances of time and place" in real-time. Any analysis that imagines otherwise, Hayek warned, simply "assumes the problem away." Recognizing that local knowledge was ubiquitous and indispensable led Hayek to conclude that efficient choices must be decentralized.

But the rise of powerful machine learning systems – the kind of transformative AI that is the subject of this volume – demands a reappraisal of Hayek's assumptions and thus their implications for organizational design at the scale of firms, sectors, and indeed entire economies. By "transformative AI" we mean systems whose aggregate capability and speed would approximate, in Dario Amodei's phrase, "a country of geniuses in a datacenter," i.e., AI that can perform all cognitive work that humans can perform and thereby reconfigure possibilities for production and coordination (Amodei, 2024). The case for taking this possibility seriously is bolstered by regularities in scaling and in inputs: empirically, model performance obeys power-law scaling with increases in data, compute and the number of parameters (Kaplan et al., 2020; Hestness et al., 2017), with revised "compute-optimal" training rules further improving efficiency (Hoffmann et al., 2022). In parallel, training compute at the frontier has grown on the order of 4–5× per year since 2010, reflecting sustained increases in spending and infrastructure (Sevilla and Roldán, 2024); industry tracking likewise documents rapid capability gains across benchmarks (Stanford HAI, 2025; Maslej et al., 2025). Taken together, these trends suggest continued movement toward significantly more powerful AI systems.

In this chapter, we investigate the implications of the rise of AI for the organization of economic production, starting from the observation that many claims about the inherently dispersed nature of knowledge and information processing are now in question. We offer a simple conceptual framework to understand the potential of transformative AI to centralize decision-making, and early evidence of the causes and consequences of this centralization.

Our analysis is based on a straightforward application of the property rights theory of the firm (Grossman and Hart, 1986; Hart and Moore, 1990), which we outline in section 2. In section 3, we highlight one of the strengths of centralized planning: it can take into account the myriad interdependencies that inevitably exist among local decisions. However, the knowledge needed to address these interdependencies is dispersed, a fact which we introduce into our analysis in section 4. This analysis mirrors Hayek's argument that it can be infeasible to centralize this knowledge when most of it is inalienable, and so decentralization is necessary.

In section 5 we capture distinct argument for decentralized decision-making noted by (Simon, 1955): human information processing capacity is bounded. Even if the entirety of the relevant information in a medium-sized business, let alone an industry or whole economy, could be transmitted to a central decision-maker, no human brain has the capacity to consider and analyze it all. We model the rise of AI as affecting the allocation of decision rights through these two channels: (i) changes in the alienability of information and (ii) increases in information processing capacity. We discuss several countervailing forces that could instead push toward decentralization, even in a world with transformative AI in section 6.

We are not the first to point out AI's potential for greater centralization of decision-making. As we discuss in more detail in section 7, there is a growing empirical literature documenting increasing concentration across industries (Autor et al., 2020; De Loecker et al., 2020) and some research linking this phenomenon to the rise of information technologies (Brynjolfsson et al., 2008; Bessen, 2020; Brynjolfsson et al., 2023). However, these analyses have not yet incorporated the change that could result from transformative AI. Other writers have framed the observation that machine learning may centralize economic power in political rather than strictly economic terms, and we too discuss the political implications of centralization in section 8.[1] A few have offered more casual and speculative explorations of how AI enables centralization (Bastani, 2019; Drago and Laine, 2025) or offered policy-relevant analyses that share the spirit of this chapter (Brynjolfsson and Ng, 2021; Agrawal et al., 2022; Acemoglu, 2023).

2 Framework

We use a three-party incomplete contracts framework (Grossman and Hart, 1986; Hart and Moore, 1990) to analyze the economic efficiency of regimes in which decision-making is centralized and decentralized.

The parties $N = \{M_1, M_2, H\}$ include two local "managers," M_1 and M_2, and a central "headquarters," H. There is a set of assets $a \in A$. For example, H could be the central

[1] Some have argued – noting the ironies – that the very corporations held up as symbols of modern capitalism also reveal new infrastructural possibilities for socialism in the 21st century (Jameson, 2016; Phillips and Rozworski, 2019; Morozov, 2019).

office of a cafe franchise and M_1 and M_2 could be local managers of cafes, or H could be a dispatcher and M_1 and M_2 could be field service technicians.

An ownership regime $\rho\colon N \to 2^A$ is a map assigning to each party the subset of assets it owns. Assets are excludable,[2] and for a coalition $S \subseteq N, \rho(S) = \cup_{k\in S}\rho(k)$. Throughout our analyses, we will consider two ownership regimes: a "centralized" regime in which the HQ owns as many assets as is technologically feasible, and a "decentralized" regime in which the local managers own as many assets as is technologically feasible.

Contracts are incomplete. That is, there are at least some contingencies that the parties cannot foresee and as a result, investments are non-contractible. At stage 0, the parties (cooperatively) choose an ownership regime. At stage 1, the parties make non-contractible relationship-specific investments x_k and incur convex costs $c_k(x_k)$, for $k = 1, 2, H$. These investments could represent mastering a popular neighborhood cuisine and supply network in the case of the local cafe, or building a reputation for expertise and reliability in the case of the field service operation. At stage 2, output V is realized and the parties divide the surplus via the Shapley value, where each hypothetical "coalition" can deploy the assets controlled by its members.

More specifically, the output or production surplus of a coalition $S \subseteq N$ is given by $V(S, \rho(S), x)$ where $x = (x_1, x_2, x_H)$ is the investments made by the members of the coalition S and $\rho(S)$ is the set of assets owned by the members of the coalition. Throughout, we make the following assumptions on V: (i) it is strictly increasing in each party's investment, (ii) it exhibits diminishing returns to each party's investment, and (iii) it exhibits weak complementarities in the parties' investments. Often we will suppress the dependence of V on the investment decisions. Furthermore, we assume that manager-manager coalitions cannot form: the only coalitions that can form are singleton coalitions, coalitions with one manager and headquarters, and the grand coalition (with the two managers and headquarters).[3]

In the following sections we will make specific assumptions on the technological environment. In particular, we make assumptions about the nature of the assets in A, which will constrain both how they are deployed in production and the ownership regimes. As in the usual incomplete contracts logic, the ownership regime determines residual rights of control, affecting the value of different coalitions and therefore the parties' incentives to invest ex-ante. Throughout, the same first-order conditions for managers M_i $i = 1, 2$ and headquarters H define incentives. We wrote them below, denoting V_k to be the marginal value of k's investment for the coalition[4].

General Investment Incentives for M_i and H

$$\tfrac{1}{3}V_i(N, A) + \tfrac{1}{6}V_i(M_i, H, \rho(M_i, H)) + \tfrac{1}{2}V_i(i, \rho(i)) = c_i'(x_i) \tag{1}$$

$$\tfrac{1}{3}V_H(N, A) + \tfrac{1}{6}\sum_{i=1,2} V_H(M_i, H, \rho(M_i, H)) + \tfrac{1}{3}V_H(H, \rho(H)) = c_H'(x_H). \tag{2}$$

The different technological assumptions that we explore will often eliminate several terms in (1)-(2) and determine the relative magnitudes of others.

[2]For all distinct parties $k \neq j$, $\rho(k) \cap \rho(j) = \varnothing$ and $\cup_{k\in N}\rho(k) = A$.

[3]This restriction on coalitions follows the star communication network restriction studied in Myerson (1977).

[4]That is, the marginal product of x_i given coalition S and investment profile x is

$$V_i(S, \rho(s)) \equiv \frac{\partial V}{\partial x_i}(S, \rho(S), x).$$

3 The Need for Coordination Favors Centralization

In many production settings, distinct local assets generate value only when their use is coordinated. When decisions are technologically interdependent, a central hub can translate those interdependencies into higher surplus by choosing compatible actions and enforcing them.

The communication architecture also matters: a hub-and-spoke (HQ with each manager) replaces the $n(n-1)$ lateral links of full decentralization with only n links to the center, making it cheaper to aggregate information and propagate consistent plans (Malone and Smith, 1988). This helps not only communication, but also incentives and enforcement. Instead of n contracts with the hub, they might they require $n(n-1)$ pairwise contracts. [5] Furthermore, in a hub-and-spoke system, the hub has stronger leverage for enforcing contracts than any individual manager. As the nexus of contracts, it has a variety of other incentive tools, including withholding access to the whole network, while in pairwise contact, individual managers can only withhold their own assets. (Jensen and Meckling, 1992)

As a result, local units may not internalize system-wide complementarities absent central ownership and authority. We capture these interdependencies with two reduced-form primitives: a coordination multiplier $\lambda \geq 1$ that scales the marginal product of a manager's investment when assets are jointly directed by HQ, and a local autonomy parameter $\alpha \in [0,1]$ that measures how much of that marginal product the manager can realize alone.

Assumption 1 (Interdependence of assets)**.** *For each manager $i = 1,2$, there is a coordination multiplier $\lambda \geq 1$ such that for any coalition S with $k \in S$, $k = 1, 2, H$, and for any asset $a_j \in A$,*

$$V_k(S, a_1, a_2) = \lambda V_k(S, a_j),$$

and a local autonomy multiplier $\alpha \in [0,1]$

$$\alpha \equiv \frac{V_i(M_i, a_i)}{V_i(M_i, H, a_i)},$$

with all derivatives evaluated at the relevant equilibrium investment level.

We now consider the investment decisions in the two regimes. In the centralized ownership regime,[6] investment decisions are determined by (3) and (4).

Centralized Ownership

$$\tfrac{1}{3}V_i(N) + \tfrac{1}{6}\underbrace{V_i(M_i, H, a_1, a_2)}_{=\lambda V_i(M_i, H, a_i)} = c_i'(x_i) \tag{3}$$

$$\tfrac{1}{3}V_H(N) + \tfrac{1}{6}\sum_{i=1,2} V_H(M_i, H, a_1, a_2) + \tfrac{1}{3}V_H(H, a_1, a_2) = c_H'(x_H). \tag{4}$$

In the decentralized ownership regime,[7] investment decisions are determined by (5) and (6).

Decentralized Ownership

[5] The number of contracts is far larger if contracts between every coalition are considered, approximatly 2^{2n-1} for large n

[6] In the centralized regime, $\rho^{\text{cen}}(H) = \{a_1, a_2\}$ and for $i = 1, 2$, $\rho^{\text{cen}}(M_i) = \varnothing$.

[7] In the decentralized regime, $\rho^{\text{dec}}(H) = \varnothing$ and for $i = 1, 2$, $\rho^{\text{dec}}(M_i) = a_i$.

$$\tfrac{1}{3}V_i(N) + \tfrac{1}{6}V_i(M_i, H, a_i) + \tfrac{1}{2}\underbrace{V_i(M_i, a_i)}_{=\alpha V_i(M_i,H,a_i)} = c_i'(x_i) \tag{5}$$

$$\tfrac{1}{3}V_H(N) + \tfrac{1}{6}\sum_{i=1,2} V_H(M_i, H, a_i) = c_H'(x_H). \tag{6}$$

We can understand the efficiency of each ownership arrangement by comparing the first-order conditions. For headquarters, comparing (4) and (6) shows that incentives are unambiguously stronger under centralization: HQ enjoys the extra singleton term and, by Assumption 1, the pair terms are weakly larger (strictly if $\lambda > 1$). For managers, centralization strengthens the pair term by a factor λ whereas decentralization adds a singleton fallback term weighted by $\frac{1}{2}\alpha$. Thus these investment incentives readily give the following sufficient condition for centralization to be more efficient than decentralization.

Proposition 1. *If $\lambda > 1 + 3\alpha$,[8] then centralized ownership yields strictly higher surplus than decentralized ownership.*

Note that in the case of zero local autonomy, $\alpha = 0$, any positive coordination gain ($\lambda > 1$) will make centralization more efficient. More generally, the more interdependencies matter (α low, λ high) the more attractive centralization becomes. Because this is a sufficient (not necessary) condition, centralization can still dominate even when $\lambda \leq 1 + 3\alpha$, for example if the marginal value of headquarters' investment is particularly large relative to that of the managers.

4 Local Knowledge and the Alienability of Information Assets

The previous section showed that the need for coordination leads to more centralized decision-making and asset ownership. However, Hayek argued that local knowledge makes such centralization infeasible, and so decisions should still be made by the person on the spot who holds the knowledge and information most relevant to the decision. In this section we formalize Hayek's epistemic argument by introducing the concept of "local knowledge" into the framework outlined in section 2.

For each local manager i there are two assets: an information asset ("local knowledge") $a_{I(i)}$ initially possessed by the corresponding manager, and a tangible asset a_i. We assume that each information asset is strictly complementary with its corresponding tangible asset—that is, in order for a_i to be valuable in a coalition, $a_{I(i)}$ must also be controlled by the coalition.[9]

4.1 The Need for Local Knowledge Favors Decentralization

We assume first, to capture Hayek's argument, that the information assets are *inalienable*, i.e. the only party that can own information asset $a_{I(i)}$ is the local manager M_i. This assumption determines the possible centralized and decentralized ownership regimes. Recall, we take the "centralized" regime to be where headquarters owns as many assets as is technologically feasible, so here,

$$\rho^{\text{cen}}(H) = \{a_1, a_2\} \quad \text{and} \quad \rho^{\text{cen}}(M_i) = \{a_{I(i)}\} \text{ for } i = 1, 2,$$

[8] Note, α in this condition is evaluated at the decentralized equilibrium investment level.

[9] Formally, for $i = 1, 2$, if $a_{I(i)} \notin \rho(S)$ but $a_i \in \rho(S)$, then $V(S, \rho(S)) = V(S, \rho(S) \setminus a_i)$. This set up mirrors (Brynjolfsson, 1994).

and the decentralized regime is one in which the local managers own their assets, so here,

$$\rho^{\text{dec}}(H) = \varnothing \quad \text{and} \quad \rho^{\text{dec}}(M_i) = \{a_i, a_{I(i)}\} \text{ for } i = 1, 2.$$

The investment incentives in the centralized regime are given by (7) and (8).

Centralized Ownership Under Inalienable Information Assets

$$\tfrac{1}{3}V_i(N, A) + \tfrac{1}{6}\underbrace{V_i(M_i, H, a_1, a_2, a_{I(i)})}_{=V_i(M_i, H, a_i, a_{I(i)})} = c_i'(x_i) \tag{7}$$

$$\tfrac{1}{3}V_H(N, A) + \tfrac{1}{6}\sum_{i=1,2}\underbrace{V_H(M_i, H, a_1, a_2, a_{I(i)})}_{=V_H(M_i, H, a_i, a_{I(i)})} = c_H'(x_H). \tag{8}$$

Meanwhile, under the decentralized ownership regime, investment incentives are given by (9) and (10).

Decentralized Ownership Under Inalienable Information Assets

$$\tfrac{1}{3}V_i(N, A) + \tfrac{1}{6}V_i(M_i, H, a_i, a_{I(i)}) + \tfrac{1}{2}V_i(M_i, a_i, a_{I(i)}) = c_i'(x_i) \tag{9}$$

$$\tfrac{1}{3}V_H(N, A) + \tfrac{1}{6}\sum_{i=1,2} V_H(M_i, H, a_i, a_{I(i)}) = c_H'(x_H). \tag{10}$$

These two sets of first-order conditions show that giving ownership of the tangible asset to the managers increases the incentives of the managers while leaving the headquarters' incentives unchanged.

To see this, first note that in the terms involving a coalition of headquarters and just one manager i, there are three assets present: the manager i's information asset $a_{I(i)}$ (controlled by i) and the two tangible assets a_1, a_2 controlled by H. However, the presence of a_j does not affect production without its complement $a_{I(j)}$, so the second term in (8) is in fact the same as the second term in (10), and thus headquarters' incentives are exactly the same in the two regimes.

Meanwhile, comparing (7) to (9), we see that managers' incentives are strictly stronger given the extra third term that appears in (9) but not (7), which represents what the manager can obtain in a coalition on their own with control of the tangible asset that complements their information. This comparison of incentives is unambiguous and directly translates to an increase in surplus under the decentralized regime.

Proposition 2. *If the information assets $a_{I(i)}$ are inalienable, then joint surplus is higher when the managers own the complementary tangible assets a_i than it is when the headquarters owns the assets.*

This result formalizes Hayek's intuition that effective control should rest with the agent who possesses non-transferable, decision-relevant knowledge. Without this knowledge, the interdependencies that pushed toward centralization in section 3 cannot be realized. When $a_{I(i)}$ cannot be transferred, as Hayek posits is the case for local knowledge, giving the local party (M_i) decision rights over the complementary physical asset maximizes investment incentives.

In the cafe example, if the local manager owns the store and ovens, any extra effort she puts into learning neighborhood tastes raises not only joint revenue but also her fallback income should negotiations with the franchiser fail. When the ex-post division of surplus takes into account this possibility via the singleton coalition in the Shapley Value, the manager's ex-ante incentives for investment, and thus total surplus, are greater. Under franchiser ownership, that fallback is zero, so the manager has diminished incentives to invest in the local adaptations that could drive profits.

4.2 AI Increases Codifiability of Knowledge

Transformative AI sharply expands what counts as codifiable – and therefore transferable – "local knowledge," in three main ways: it makes *explicit knowledge* more accessible to decision-makers, it increasingly extracts *tacit know-how* once embedded in human perception and practice, and it generates *machine-native knowledge*—patterns no human could feasibly enumerate.

Explicit knowledge is knowledge that is easy to put into words or write down. AI systems can both digitize explicit knowledge at lower cost than ever before, and also absorb remarkable amounts of it into their knowledge bases. Written archives are just one – nicely literal – example of explicit knowledge, and optical character recognition techniques were one of the first practical applications of neural networks (LeCun et al., 1998). They have enabled everything from the rapid digitization of enduring texts like the Vatican Apostolic Archive, handwritten in medieval Latin that many human readers struggled to understand (Firmani et al., 2018) to the reading of fast-passing license plate numbers for use in traffic management and enforcement (Lubna et al., 2021) and the seamless uploading of physician notes into electronic health records (Hsu et al., 2022). Facts and certain forms of scientific knowledge are also explicit – questions such as "Which city is the capital of Burkina Faso?", or "How does the Haber-Bosch process synthesize ammonia?" – and increasingly encoded in single AI systems, as frontier model performance on broad factual and scientific question-answer benchmarks like MMLU (Hendrycks et al., 2021), Natural Questions (Kwiatkowski et al., 2019), TriviaQA (Joshi et al., 2017), and the biomedical PubMedQA (Jin et al., 2019) demonstrates. Physical facts about the world are another form of explicit knowledge, which AI helps to capture with unprecedented precision – for example, BMW uses AI to generate detailed 3D digital twins of factory assets used in simulation-based optimization (BMW Group, 2022).

Explicit knowledge stands in contrast to tacit knowledge (Polanyi, 1966), which is intuitively known – it's the knowledge that is easy to act on but hard to put into words. Canonical examples of tacit knowledge include driving a car (not something one can learn solely from reading a book), language learning (it's far easier to learn a language than to explicitly explain its many grammatical rules and exceptions to grammatical rules, elaborate its vocabulary alongside qualifications about common usage), and facial recognition (we intuitively recognize faces but struggle to articulate what in a face makes it distinctive). AI systems have largely mastered the latter two forms of tacit knowledge in the last decade and are making progress on the first one. Today's frontier large-language models are proficient in 50 languages[10] (OpenAI, 2024) – gaps in coverage reflect the uneven availability of training data rather than any fundamental barrier (Pava et al., 2025) – and state-of-the-art face-recognition systems, like Google's FaceNet, surpass 99.6 percent accuracy on the Labeled Faces in the Wild benchmark, outpacing human performance (Schroff et al., 2015). By ingesting billions of miles of sensor data—camera feeds, LiDAR, radar, vehicle telemetry—AI now powers a fleet of 1500 self-driving taxis (Waymos) completing over 250,000 rides a week

[10] AI-based translation tools support even more languages—Google Translate currently supports 243 languages (Peters, 2024).

in San Francisco, Austin, Phoenix and Los Angeles (Shepardson, 2025).

Beyond everyday skills, much of workplace know-how is also tacit –rooted in the heuristics developed on the job – and AI is now beginning to capture and redistribute that, too. Tools like MedGaze record thousands of eye-tracking sessions from practicing radiologists – logging where and how long experts fixate while reading X-rays – to train AI diagnostic assistants (Awasthi et al., 2025). Training on millions of transcripts allowed a large language model to capture some of the best call center agents' expertise and make it available to newer and less skilled agents (Brynjolfsson et al., 2025). In each case, AI helps to overcome Polanyi's paradox (Autor, 2014) – we now no longer need to articulate a tacit knowledge or skill for it to be automated. Meanwhile, a new marketplace now monetizes tacit expertise for AI training: companies like Mercor connect specialists to AI firms to create bespoke data – at the time of writing, their opportunities page calls for all manner of experts, from dermatologists ($270/hour) to plant experts at $30–60/hour (Mercor, 2025).

Beyond explicit knowledge and tacit knowledge, which have in the past been housed by human minds, machine learning has enabled the discovery or creation of entirely new forms of knowledge that no human mind – or set of human minds – could comprehend. From fraud detection tools (Dal Pozzolo et al., 2018) and protein-folding predictions that suggest novel drug targets (Jumper et al., 2021; Evans et al., 2022), to anomaly-detection systems that allocate firefighting resources preemptively (Jain et al., 2020) and high-frequency trading algorithms that exploit microsecond price discrepancies (Budish et al., 2015) — AI continuously pushes the frontier of what counts as codifiable, economically relevant knowledge.

Collectively, these three developments extend the codifiable frontier at an accelerating pace. A growing share of economically relevant facts, heuristics, and predictive signals now resides in databases, embeddings, and model weights—readily stored, duplicated, and recombined at negligible cost.

We offer more detail about these developments in Appendix ??.

4.3 AI Tips the Scales Toward Centralization

When transformative AI makes the local information asset codifiable and transferable, we can treat it as alienable. Using strict complementarity, we can bundle the information and tangible assets so that the coordination logic from section 3 applies verbatim.

For each manager i, the information asset $a_{I(i)}$ and the tangible asset a_i are strict complements in the sense that a coalition's marginal product in x_i is zero unless *both* a_i and $a_{I(i)}$ are available to that coalition. If $a_{I(i)}$ is codifiable and transferable, we can bundle

$$\tilde{a}_i \;\equiv\; \{\, a_i,\; a_{I(i)} \,\}$$

and restrict attention to ownership regimes that keep complements together.

We consider the same two regimes as before, now phrased in terms of the bundled assets $\tilde{A} = \{\tilde{a}_1, \tilde{a}_2\}$: the centralized regime[11] and the decentralized regime.[12]

Because each $\tilde{a}_i$ moves as a unit, the coalition values with bundles are the same functions as in section 3, and the relevant first-order conditions revert exactly to (3)–(4) under centralization and (5)–(6) under decentralization (with a_i understood to denote $\tilde{a}_i$). Therefore, the same sufficient condition from Proposition 1 applies here.

[11] In which $\rho^{\text{cen}}(H) = \tilde{A}$ and $\rho^{\text{cen}}(M_i) = \varnothing$ for $i = 1, 2$.

[12] In which $\rho^{\text{dec}}(M_i) = \{\tilde{a}_i\}$ for $i = 1, 2$ and $\rho^{\text{dec}}(H) = \varnothing$.

Intuitively, once information can move, *co-locating* decisions with information can be achieved by transferring the information to H together with its physical complement, restoring the coordination gains characterized in section 3. This aligns with insights in Jensen and Meckling (1992) and Brynjolfsson and Mendelson (1993), and challenges Hayek's intuition when he wrote, for example, "Practically every individual has some advantage over all others because he possesses unique information of which beneficial use might be made, but of which use can be made only if the decisions depending on it are left to him."

If some of HQ's own information assets also become codifiable and transferable to the periphery, the logic can run the other way: bundling those with local physical assets can make decentralization feasible where it previously was not. Nonetheless, when interdependencies are significant (α low and λ high), the coordination advantages in section 3 will still make centralization more attractive.

5 Information Processing Capacity

Hayek's side of the socialist calculation debate pointed not just to the difficulty of codifying knowledge, but also to the finite processing power of any single decision-making entity. "The problem," he wrote, is "how to extend the span of our utilization of resources beyond the span of the control of any one mind" (Hayek, 1945). This challenge, rooted in the inherent limits of human cognition, finds a more formal expression in Herbert Simon's concept of "bounded rationality" (Simon, 1955). Because attention is scarce, modern organizational economics treats decision-making as necessarily distributed across agents, teams and hierarchies rather than vested in a single planner (Arrow, 1974; Radner, 1993; Radner and Van Zandt, 1992; Bolton and Dewatripont, 1994; Garicano, 2000). Here we first discuss how bounded information processing has historically hampered the value of centralized control in production, and then present evidence that AI is rapidly expanding the amount of information that can be processed by a single entity.

5.1 Limits on Information Processing Push Toward Decentralization

A single mind, or collection of minds in a firm, can only process so much information. While the brain's unconscious processing capacity is vast, estimates for conscious, deliberate reasoning are as low as 60 bits per second (Csikszentmihalyi, 1990). This biological bottleneck is more or less fixed. While estimates vary, some place the computational power of the human brain around $10^{1}5$ FLOP/s, a threshold that modern supercomputers can now surpass (Carlsmith, 2020). And moreover, not all the information contained in the minds of the people who make up a firm can be deployed all at once. There is friction – information held by different people cannot be instantaneously accessed, introducing delays and misunderstandings, imprecision due to incentive misalignment, and other communication challenges. As a result, organizations pay processing costs each time knowledge crosses a human boundary – this puts an overall limit on the amount of information any collection of human minds can effectively process.

These bounds mean that even if centralized control may be more efficient and information is alienable, there may be other bottlenecks to centralization because of the sheer amount of information that needs to be processed to realize the valuable interdependencies.

To formalize how bounds on information processing limit the returns to centralized ownership, we extend the GHM environment to N local managers and a central headquarters on a star communication network as in (Myerson, 1977). Let the parties be

$$N = \{M_1, \dots, M_N, H\}, \qquad A = \{a_1, \dots, a_N\},$$

where a_i is manager i's local asset bundle. As before, for a coalition $S \subseteq N$, the assets available to S under ρ are $\rho(S) := \{a \in A : \rho^{-1}(a) \in S\}$, and its surplus is $V(S, \rho(S), x)$ with $x = (x_{M_1}, \dots, x_{M_N}, x_H)$.

As before, we compare two regimes. The *centralization* regime ρ^{cen} with $\rho^{\text{cen}}(H) = A$ and *decentralization* ρ^{dec} with $\rho^{\text{dec}}(M_i) = a_i$ for all i.

Assumption 2 (Limited information processing)**.** *Headquarters can* process *at most* $\bar{K} \in \{0, 1, \dots, N\}$ *assets at stage 2. Let* $s_i \in \{0, 1\}$ *indicate whether asset i is processed, with* $\sum_{i=1}^{N} s_i \leq \bar{K}$*. Interdependence gains are summarized by* $\lambda \geq 1$ *and realized only for processed assets under centralization.*

It is useful to introduce notation for the marginal value of each manager's investment in a bilateral coalition with HQ when only one asset $a_i = \rho(M_i)$ is present: $\tilde{V}_i = V_i(M_i, H, \rho(M_i))$ for all $i = 1, \dots, N$. Similarly, we define the marginal value of H's investment in a bilateral coalition with M_i when only asset i is present by $\tilde{V}_H(i) = V_H(M_i, H, a_i)$. With this notation, we can extend the interdependence assumption (Assumption 1) to the N manager case.

Assumption 3 (Interdependence of assets; N managers)**.** *Assume that for any bilateral coalition between a manager* M_i *and headquarters* H*, with assets* $\rho(M_i), \rho(H)$*,*

$$V_i(M_i, H, \rho(M_i), \rho(H)) = \lambda \tilde{V}_i \qquad V_H(M_i, H, \rho(M_i), \rho(H)) = \lambda \tilde{V}_H(i).$$

Note that the symmetry imposed by this assumption is quite strong—anytime an asset a_i is coordinated with any additional assets, the marginal value of investment is multiplied by the same factor λ, regardless of *which* other assets are present or *how many* there are.[13]

Under centralization with capacity $\bar{K}$, incentives are determined by the following first order conditions. Note that the conditions from section 2 must be adapted to account for N players using the Shapley Value on a star coordination network; for ease of exposition we write the coefficient on pairwise coalitions as $\Phi(N) \equiv \frac{N-1}{2(N+1)}$.

Centralized Ownership with Processing Limits

$$\tfrac{1}{N+1} V_i(N, A) + \Phi(N) s_i \lambda \tilde{V}_i = c_i'(x_i) \tag{11}$$

$$\tfrac{1}{N+1} V_H(N, A) + \Phi(N) \sum_{i=1}^{N} s_i \lambda \tilde{V}_H(i) + \tfrac{1}{N+1} V_H(H, A) = c_H'(x_H) \tag{12}$$

Under decentralization with capacity $\bar{K}$, incentives are determined by

Decentralized Ownership with Processing Limits

$$\tfrac{1}{N+1} V_i(N, A) + \Phi(N) \tilde{V}_i + \tfrac{1}{2} \alpha_i \tilde{V}_i = c_i'(x_i). \tag{13}$$

$$\tfrac{1}{N+1} V_H(N, A) + \Phi(N) \sum_{i=1}^{N} \tilde{V}_H(i) = c_H'(x_H) \tag{14}$$

[13] In other words, the coordination multiplier is uniform across all pairwise matchings (a_i, a_j) and also across all larger combinations (a_i with any subset of the remaining assets). This simplifies the exposition, and can be read conservatively by treating λ as a lower bound on coordination gains across configurations. A richer version could let the multiplier depend on the number or identity of complements.

The grand-coalition terms (the first terms on the LHS of (11)–(14)) are identical across regimes because the grand coalition controls all assets and includes H either way so they are irrelevant to regime comparisons. By Shapley weighting on the star, the "pair-with-H" terms receive weight $\Phi(N)$, while the manager's singleton fallback has coefficient 1/2. Under centralization, only the $\bar{K}$ assets that headquarters actually processes ($s_i = 1$) realize the coordination gain $\lambda \tilde{V}_i$ while under decentralization, each manager gets the unscaled pair term $\tilde{V}_i$ plus a singleton fallback contribution $\frac{1}{2}\alpha_i \tilde{V}_i$.[14] By a similar logic to that of section 3, with an additional symmetry condition, these comparisons readily give a simple sufficient condition for the efficiency of centralized control.[15]

Proposition 3. *Assume symmetry across managers, i.e. for all $i \neq j, \tilde{V}_i = \tilde{V}_j$. If*

$$\lambda > \frac{N}{\bar{K}}\Big(1 + \frac{N+1}{N-1}\,\alpha\Big),$$

then centralized ownership yields higher surplus than decentralization.

Intuitively, the right-hand side of the sufficient condition rises with N and with local autonomy α (which strengthens decentralized fallbacks) and falls with processing capacity $\bar{K}$. When $\bar{K} = N$ the condition simplifies to $\lambda > 1 + \frac{N+1}{N-1}\alpha$ (recovering $\lambda > 1 + 3\alpha$ when $N = 2$). As $\bar{K}$ increases—for example via AI-enabled processing—the sufficient condition for centralization becomes easier to satisfy.

5.2 *AI Eases Limits on Information Processing Capacity*

In the past, the headquarters' effective processing capacity $\bar{K}$ was limited because coordinating interdependent local assets means taking in, storing, and reasoning over large, fast-changing information within each decision cycle. Recent AI advances relax this bottleneck in four main ways.

The first is larger working memory. Modern systems can read and condition on much longer inputs in one go, often in the range of two hundred thousand to one million tokens, where a token is roughly a short word or part of a word. That means a single run can take in entire databases at once rather than juggling many small snippets (Anthropic, 2024; Pichai and Hassabis, 2024; OpenAI, 2025).[16]

The second channel is higher inference throughput. Software improvements make each decision faster and cheaper in token terms, which raises how much information can be processed per planning cycle. FlashAttention reorganizes how attention moves data so it runs faster and uses less memory (Dao et al., 2022). PagedAttention, as used in vLLM, manages memory so that many long requests can be served efficiently (Kwon et al., 2023). Speculative decoding speeds up generation by drafting likely continuations and confirming them, which cuts waiting time without changing the final answer (Leviathan et al., 2023).

[14] Note that the coefficient on the managers' standalone value is $\frac{1}{2}$ regardless of the size of N. To see why, recall the random-order view of the Shapley value in which we average each player's marginal contribution over all arrival sequences. For a manager M_i on a star, any sequence where the hub H arrives after M_i yields a standalone marginal $\tilde{V}_i$, and this event has probability $\frac{1}{2}$ independent of N.

[15] Note that here we cannot compare the incentive conditions manager-by-manager comparisons, instead we must sum across the managers. Because payoffs use Shapley weights, each manager's first order condition is linear in marginal contributions, and with separable convex costs the sum of these first-order conditions equals the derivative of a common potential (expected surplus). Summing therefore provides a valid aggregate incentive comparison.

[16] That said, today's models still struggle to perfectly use long context windows, so longer windows work best together with methods that help the model focus (Liu et al., 2024).

The third channel is external memory and tool use. Long context is not the only way to scale what a decision center can consider. Retrieval methods let the model fetch only the most relevant pieces from a much larger store, so the system need not load everything at once. RETRO shows how retrieval can be built into training so that a model can draw on trillions of tokens without growing its internal size (Borgeaud et al., 2022). At inference time, vector search systems such as FAISS pull the nearest documents in milliseconds even when the index is very large (Johnson et al., 2017). Models can also call tools when needed, such as search, code, or databases, which moves parts of the task outside the model while keeping the reasoning loop intact (Schick et al., 2023; Yao et al., 2023). In practice this makes memory and compute elastic, since the center pulls in only what a particular decision requires.

The fourth channel is algorithmic routing that concentrates compute where it matters most. Mixture-of-Experts models switch on only a few expert modules for each token, so the model's total capacity can grow without growing the cost of every step (Fedus et al., 2021). Long-document readers like Longformer and Performer reduce how much of a text the model has to look at, which makes very long inputs affordable while preserving the ability to spot the important parts (Beltagy et al., 2020; Choromanski et al., 2021). New sequence models like Mamba maintain a compact running state as they process long streams, which also keeps cost roughly linear in length (Gu and Dao, 2023). These ideas allow the system to devote more effort to the assets and interactions that actually drive coordination gains.

5.3 AI Again Tips the Scales Toward Centralization

Returning to the key condition in Fact 3, we see that as $\bar{K}$ increases, the sufficient condition for the domination of centralization becomes easier to satisfy. This is because headquarters can realize the coordination synergies across a larger and larger fraction of the firm's assets. While decentralized ownership still provides stronger individual investment incentives for managers, the value created by this effect becomes increasingly outweighed by the surplus generated through firm-wide coordination.

6 Countervailing Forces

While we've laid out some reasons that powerful AI systems strengthen the economic efficiency of centralized decision-making, there are three counterarguments to consider.

First, AI may simply not be capable of making all types of economic decisions or processing all economically relevant information. Humans could retain an advantage in some categories and this may, in turn, necessitate decentralization for those types of decisions. Second, even if we assume that AI is capable of all decisions and there are no reasons that humans need to be involved in decision-making, there may still be a decentralization of decision-making among different AI entities. Third, even if AI is capable of making economic decisions as well as humans, we may still decide that we want humans to make certain decisions. This may be expressed through preferences or legislation that prevents centralization. We consider each of these cases in turn.

6.1 Limits to codifiability or machine information processing

As Hayek emphasized, not all information is readily codifiable. Humans still hold advantages in embodied and affective skills, e.g. perception, dexterity, and social sensing. Competent people can still fold laundry, button a shirt, or throw a curve ball more reliably than machines, and they often better read micro-expressions and vocal cues. When decisions hinge on such locally perceived signals, delegating to the human on the spot will continue to dominate centralized control.

A second limit to the codifiability of knowledge comes from the "long tail." Many domains exhibit long-tailed distributions: a mass of common patterns but also a vast set of rare situations with little or no data. ML systems trained on historical data handle the frequently observed mass but can fail in the tail. For example, large language models struggle to learn long-tail knowledge (Kandpal et al., 2023), and humans retain a comparative advantage on rare diagnostic cases (Agarwal et al., 2025). Moreover, people often generalize better than machines from few examples—classic one-shot learning (Fei-Fei et al., 2006; Lake et al., 2015). A toddler may recognize elephants after a couple of pictures; current systems typically need many. In medicine, radiologists "know the long tail" (Langlotz, 2019), which is precisely where safety-critical errors matter most. Furthermore, human preferences themselves are a form of knowledge that machines may always struggle to learn perfectly. While many digital platforms run on predictive models of consumer preferences,[17] these predictive models work less well when preferences are highly idiosyncratic.

That said, these claims push against the premise of this Conference Volume. They assume enduring human superiority on some capabilities. Over time, AI systems may accumulate data deeper into the tail and, more importantly, improve few-shot generalization, shrinking human advantages. If so, the frontier of uniquely human value will recede—remaining real, but increasingly concentrated in rarer, more idiosyncratic cases.

6.2 When would a fully AI-powered economy be decentralized?

Even in a world where all decisions are made by AI, there may be reasons to decentralize decision-making because communication is neither instantaneous nor perfectly reliable. First, propagation is bounded (roughly at the speed of light) and links can degrade; as a result, a local decision-maker can sometimes react faster and more dependably than a distant controller when milliseconds matter. This shows up at multiple scales: at the hardware level, signal delay across a chip—on the order of picoseconds per inch—can bind at high clock speeds (Ott, 2009); in markets, algorithmic trading benefits from colocated autonomy (Steiner, 2010); and in combat, drone operations may require millisecond responses, making local autonomy tactically valuable. Second, when the distance between nodes in a communication network is especially large, decentralization may be indispensable: communications between Earth and Mars impose round-trip lags of about 4.3–21 minutes, so landers and rovers must act autonomously (NASA Jet Propulsion Laboratory, 2023). Taken together, finite latency and the realities of physics imply that even in an all-AI world, distributing decision rights toward the edge can yield operational and tactical advantages whenever required reaction times approach—or fall below—end-to-end communication delays.

Furthermore, even among machines, contracts remain incomplete: the world's contingencies cannot be fully specified ex-ante, and the combinatorics of real tasks outpace any fixed model, especially when the machines themselves multiply the number of contingencies that need to be considered, creating a "red queen" scenario. As in human organizations, non-contractible investments still matter, so distributed AI agents may need residual rights

[17] Consider, for example Amazon's anticipatory shipping (Spiegel et al., n.d.; Opam, n.d.), and the not un-common refrain in popular media that digital platforms and their algorithms "know you better than you know yourself" (Carmichael, 2014; Harari, 2017; Thompson, n.d.).

of control to act when unforeseen states arise.

That said, the direction of travel could still favor centralization if communication frictions keep falling and frontier models keep scaling. In many workloads, a single strong agent with rich tool use can outperform multi-agent schemes: recent studies find multi-agent LLM systems frequently fail to beat robust single-agent baselines and can even degrade accuracy (Cemri et al., 2025); likewise, a single search agent has outperformed multi-agent search variants (Nguyen et al., 2025).

6.3 Legislative requirements for decentralization

While the focus of this article has been on the economic efficiency of alternative arrangements for ownership and control, the "ownership regimes" of organizations at all scales—from small businesses to entire economies—are, of course, not made on the basis of economics alone. Law, policy, and social norms set the feasible—and legitimate—set of arrangements, reflecting the fact that decentralization can have benefits beyond those we model. For instance, it may facilitate greater personal autonomy and freedom.

Competition rules (merger control, structural separation) can limit centralization even when scale economies favor it. Data governance—privacy, localization, purpose limits—can block large, centralized data pools, while interoperability and data-portability mandates can open paths for decentralized entrants. Sectoral laws in finance, health, transport, and criminal justice may continue to require human oversight, contestability, and auditability; that often anchors decisions where traceability is strongest, not where coordination is cheapest. Liability and insurance rules for AIs are also yet to be written. Professional licensing (e.g. for pilots, physicians, engineers) impose human-in-the-loop obligations that could, if upheld, continue to distribute decision rights.

Even without mandates, people may continue to want humans in control of consequential or meaning-laden decisions: a judge for sentencing, a pilot during critical phases, players on the field, or a human coach, poet, or therapist. These preferences create *de facto* limits on full centralization and sustain local agency where legitimacy, dignity, or narrative matter.

6.4 Summary

To be sure, there are some important counterarguments to the idea that even AI that surpasses human cognition would lead to centralization. Some types of information might defy codification, and some types of information processing might be better done locally. Thus, not every decision would be centralized. But if transmission and processing speeds for machine-readable information are orders of magnitude larger than today, and the capabilities of AI grow commensurately, then these cases might account for a smaller and smaller share of the economy. Thus, in the absence of active countermeasures, transformative AI may lead to significantly more centralization of decision-making.

That's not to say that centralization will necessarily increase monotonically during the transition period, or that the transition period will be short. In particular, local decision-making may be advantageous during periods of turbulence, innovation or uncertainty, when rapid reactions to local information can be especially valuable. [18]

[18] For instance, there is some evidence that smaller companies have performed better after recessions (Morgan Stanley Research, 2023), and that they were especially resilient during the COVID-19 Pandemic (Wood, 2023).

7 Early Empirical Evidence of Centralization

There is already clear evidence that economic decision-making is shifting toward the center in much of the economy. Specific case studies and broader trends in market concentration both point in this direction, and many of these shifts have been explicitly linked to the growing use of information technology (IT)—transformative AI would supercharge these trends.

Some of the earliest and most striking demonstrations come from retailing. Mrs. Fields Cookies, for example, developed a headquarters-controlled expert system to prescribe store-level actions in real time in the 1980s. As company co-founder Debbi Fields put it, "We have removed the decision-making process from the store level. The manager's responsibility is to execute the plan – not to plan" (Richman, 1987). The system dictated when to mix dough, when to bake, which varieties to emphasize, whether to call in extra labor, and even when to hand out free samples, all based on live traffic forecasts (Harvard Business School, 1990). In short, "*the system not only tells her what's happening, it tells the stores what to do about it*" (Richman, 1987). By removing local discretion, Mrs. Fields used IT to optimize quality, labor and customer experience across all outlets.

Large multi-product retailers soon adopted similar playbooks – none more aggressively than Walmart, as mentioned in section 1. Beginning in 1987, Walmart linked every store to Bentonville via what was then the nation's largest private satellite network, giving headquarters real-time visibility into SKU-level sales. By 1991 it had launched Retail Link, an extranet that auto-generated store-specific replenishment orders and shared live data with suppliers. Central category managers – not local store managers – decided exactly which items each outlet would stock, in what quantities, while negotiating directly with manufacturers. The result was a coast-to-coast network that behaved like one centrally optimized warehouse (Lee, 2006; Fishman, 2006).

As digital technologies have advanced, centralized business models—franchising, corporate chains, and, more recently, private-equity-backed roll-ups—have increasingly reshaped U.S. retailing. Private-equity "add-on" deals, the hallmark of roll-up strategies, made up 43 percent of all buyouts in 2002 but almost 72 percent by 2020, and they now span countless services from optometrists to car-washes (*Serial Acquisitions and Industry Roll-ups: Background Note*, 2023). At the same time, private equity's share of all U.S. corporate equity grew from about 4 percent in 2000 to nearly 20 percent in 2021 (Institute for Local Self-Reliance, 2019).

Academic evidence points to parallel changes in market structure. Decker et al. (2020) note – drawing on prior Census studies – that in retail trade the share of sales and employment accounted for by single-unit ("mom-and-pop") establishments fell from roughly one-half to one-third between 1977 and 2007, as national big-box chains spread. Smith and Ocampo (2025) find that the geographic expansion of multi-market retailers accounts for most of the post-1990 rise in national retail concentration. These structural shifts have come at the expense of smaller firms. The Institute for Local Self-Reliance (2019) reports that retailers with fewer than 100 employees captured more than half of U.S. retail spending in 1982 but only about one-quarter by 2017. The sector's four-firm concentration ratio (C4) has risen from less than 15% in the 1970s to over 40% today (Figure 1).

Increased concentration is not confined to retail. Across the U.S. private sector, the C4 index has risen 5–8 percentage points since the late 1990s, with particularly significant increases in manufacturing, banking, telecommunications and airlines, sparking what has been called the "rise of superstar firms" (Autor et al., 2020). Stock-market data tell a similar

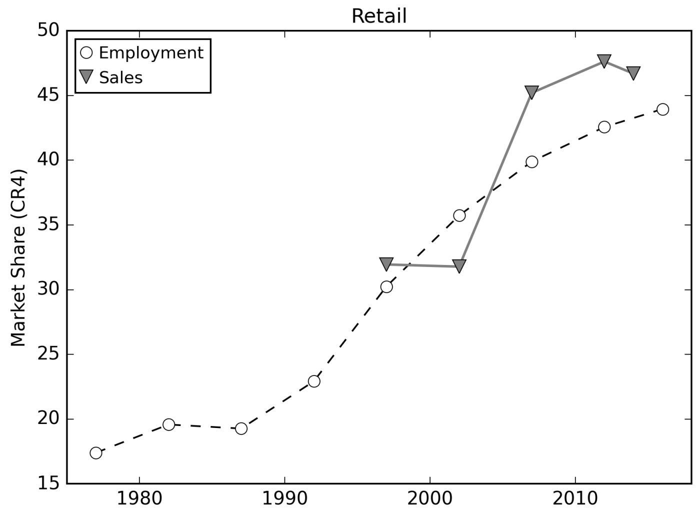

Figure 1: Growth of the U.S. retail C4 concentration ratio. Source: Brynjolfsson et al. (2023)

story: the ten largest U.S. firms now account for roughly 38% of total market capitalization – double their share in 2010 (Capitalist, 2025).

Recent literature links these patterns to IT intensity. Brynjolfsson et al. (2023) show that firms making heavier IT investments grow larger and more concentrated, primarily reflecting increases in sales not headcount consistent with earlier evidence in Brynjolfsson et al. (2008). Industry-level studies likewise find that rising concentration coincides with higher R&D and IT spending (Kwon et al., 2022), and that large software investments predict subsequent increases in market share (Bessen, 2019).

Looking ahead, AI may amplify the forces that lead to a falling share of human labor and decision-making in organizations. Sam Altman predicts the emergence of "one-person, billion-dollar" companies (Altman, 2025), and Anthropic's *Project Vend* has already let an LLM make every key decision for a small online shop—albeit with mixed results (Anthropic, 2025). While the project failed to make a profit, making some comically bad choices along the way, one can imagine a future version doing much better.

The trends to date are not necessarily predictive of what we can expect as AI becomes more powerful. However, the evidence of increasing concentration is consistent with the idea that increased codifiabilty and digitization of data, combined with increased computer processing power, makes it possible to centralize more decisions.

8 Economic Centralization and Political Power

AI can make larger, more centralized organizations economically attractive by easing information bottlenecks and enabling tighter coordination. Whether this is politically problematic will depend less on centralization itself than on how power is made accountable.

Two channels link economic centralization to the concentration of political power. First, economic concentration can increase agenda-setting and lobbying capacity, a classic prediction of political-economy models in which organized interests leverage concentrated rents

(Grossman and Helpman, 1994). Empirically, the responsiveness of US policy appears more closely aligned with economic elites and organized interests than with average citizens, consistent with concerns about unequal influence (Gilens and Page, 2014). Second, when firms that concentrate economic power also serve as information intermediaries, there are new avenues for political power. As large AI systems shape search, summarization, and content curation, they become gatekeepers of public discourse. Prior work on social media's role in misinformation and opinion formation documents the scale and stakes of such gatekeeping (Allcott and Gentzkow, 2017); adjacent scholarship highlights how data-extraction business models can amplify those dynamics (Zuboff, 2019) and how automated moderation infrastructures embed political choices (Gorwa et al., 2020). Recent analyses of generative AI underscore parallel risks for democratic representation, accountability, and trust if synthetically generated content floods civic channels (Allen and Weyl, 2024; Jain et al., 2025).

A third, longer-run channel operates through human capital and civic capacity. Education and civic skills are robustly associated with democratic stability and participation (Glaeser et al., 2007). Technological change can complement or substitute for skills in ways that alter incentives to invest in education (Acemoglu and Autor, 2011; Goldin and Katz, 2008). If AI shifts the perceived returns to certain cognitive investments, the downstream effects on civic engagement—and thus on democratic resilience—are an open and important empirical question.

Because the political consequences turn on governance, not technology alone, it is useful to note (without endorsing) strands of institutional design discussed in the literature. Work on deliberative institutions studies mechanisms for structured public input to high-stakes decisions (e.g., citizens' assemblies) (OECD, 2020; Collective Intelligence Project, 2023), and related experiments in blockchain communities include decentralized autonomous organizations (DAOs) (Hassan and De Filippi, 2021) and new forms of decentralized finance (Buterin et al., 2019). Work on data governance explores rights to access, portability, and control (GDPR, 2016; Act, 2023), collective ownership of data (Delacroix and Lawrence, 2019; Hardjono and Pentland, 2019; Cullen et al., 2025), and proposals to treat contributions of data as compensable labor (Arrieta-Ibarra et al., 2018; Posner and Weyl, 2018). Distributional proposals—including social wealth funds, or universal dividends—have also been analyzed as ways to broaden claims on AI-concentrated rents (O'Keefe et al., 2020; Huang and Manning, 2025).

9 Conclusion

This paper has examined the implications of transformative AI for the organization of economic decision-making, revisiting Hayek's epistemic critique of central planning in light of recent technological advances. We develop a simple framework in which knowledge codification and information-processing capacity are the key determinants of the optimal allocation of decision rights and ownership.

Our analysis highlights two central channels through which AI influences economic organization. First, by expanding the set of knowledge that is codifiable and transferable, AI and related technologies relax the constraints that once kept knowledge dispersed and localized, favoring decentralization. Second, by radically increasing information-processing capacity, AI enables greater integration and coordination of decisions across larger scales. Together, these channels diminish the informational and processing advantages previously held by distributed human agents, thus making centralized control more feasible and, in some contexts, more efficient.

The predictions of our theory are consistent with emerging empirical evidence on increasing concentration in U.S. industries, rising firm size and market share among IT-intensive firms, and the documented use of AI and IT to reduce discretion from local managers in retail and other industries. This suggests that, in the absence of countervailing actions, centralization is likely to intensify as AI capabilities advance.

We also note that the centralizing tendencies of AI extend beyond efficiency considerations—to potentially reshaping the distribution of economic rents and political power. As decision-making authority and residual control rights become more concentrated, human agents may see diminished bargaining power and weaker incentives to invest in human capital. If the 20th century was a century of rising human capital, the 21st may be one of increasing machine expertise. At a societal level, when people have reduced education and participation, it risks undermining their civic engagement and democratic resilience. Furthermore, concentration of economic resources and control over information flows may amplify the centralization of political power.

The framework and findings we present point to several directions for further research. One avenue is to test whether AI adoption within firms predicts shifts in centralization of decision-making authority. Another is to examine whether industries with higher AI intensity exhibit greater concentration of market power or declines in local autonomy. Further work could also investigate how AI-mediated control over information flows shapes political outcomes and public attitudes.

More broadly, our analysis suggests that the radically centralizing potential of AI may demand equally radical new ideas about how to preserve human agency and build new foundations for democracy. Understanding and addressing these questions may become one of the central economic and political issues of the coming years.

References

Acemoglu, Daron, "Would AI–Enabled Centralized Control Work?," *The International Economy*, 2023. Spring issue, pp. 49–52.

_ **and David Autor**, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4B, Elsevier, 2011, pp. 1043–1171.

Act, EU Data, "Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data (Data Act)," Official Journal of the European Union L 2023/2854 (22 December 2023) 2023.

Agarwal, Rohan, Jenna Smith, and Kevin Lee, "Human Comparative Advantage in Long-Tail Problem Solving," *Journal of Artificial Intelligence Research*, 2025, *74*, 1121–1153.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, *Power and prediction: The disruptive economics of artificial intelligence*, Harvard Business Press, 2022.

Allcott, Hunt and Matthew Gentzkow, "Social Media and Fake News in the 2016 Election," *Journal of Economic Perspectives*, 2017, *31* (2), 211–236.

Allen, Danielle and E Glen Weyl, "The real dangers of generative AI," *Journal of Democracy*, 2024, *35* (1), 147–162.

Altman, Sam, "On One-Person, Billion-Dollar Companies," Lean AI Leaderboard interview, retrieved June 2025 2025.

Amodei, Dario, "Machines of Loving Grace: How AI Could Transform the World for the Better," Essay October 2024. Includes the formulation "a country of geniuses in a datacenter".

Anthropic, "Claude 3.5 Sonnet," `https://www.anthropic.com/news/claude-3-5-sonnet` June 2024. Accessed: 2025-09-29.

Anthropic, "Project Vend: Autonomous Retail Experiment with Claude Sonnet 3.7," Technical blog post 2025.

Arrieta-Ibarra, Imanol, Leonard Goff, Diego Jiménez-Hernández, Jaron Lanier, and E. Glen Weyl, "Should We Treat Data as Labor? Moving Beyond "Free"," *AEA Papers and Proceedings*, 2018, *108*, 38–42.

Arrow, Kenneth J., *The Limits of Organization*, New York: W. W. Norton, 1974.

Autor, David, "Polanyi's Paradox and the Shape of Employment Growth," Working Paper 20485, National Bureau of Economic Research September 2014. Prepared for the Federal Reserve Bank of Kansas City Economic Policy Symposium, Jackson Hole.

_ **, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen**, "The Fall of the Labor Share and the Rise of Superstar Firms," *Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Awasthi, Akash, Anh Mai Vu, Ngan Le, Zhigang Deng, Supratik Maulik, Rishi Agrawal, Carol C. Wu, and Hien Van Nguyen, "Modeling radiologists' cognitive processes using a digital gaze twin to enhance radiology training," *Scientific Reports*, April 2025, *15* (13685).

Bastani, Aaron, *Fully Automated Luxury Communism: A Manifesto*, Verso Books, 2019.

Beltagy, Iz, Matthew E. Peters, and Arman Cohan, "Longformer: The Long-Document Transformer," *arXiv preprint arXiv:2004.05150*, 2020.

Bessen, James, "Information Technology and Increased Market Concentration," *Research Policy*, 2019, *48* (8), 103788.

Bessen, James E., "Industry Concentration and Information Technology," *Journal of Law and Economics*, 2020, *63* (3), 531–555.

BMW Group, "BMW Group publishes SORDI, the largest open-source dataset by far for super-efficient AI applications in production," `https://press.bmwgroup.com/global/article/detail/T0375993EN/bmw-group-publishes-sordi-the-largest-open-source-dataset-by-far-for-super-efficient-ai-appli` March 2022. Press release, PressClub Global.

Bolton, Patrick and Mathias Dewatripont, "The Firm as a Communication Network," *The Quarterly Journal of Economics*, 1994, *109* (4), 809–839.

Borgeaud, Sebastian, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack W. Rae, Erich Elsen, and Laurent Sifre, "Improving Language Models by Retrieving from Trillions of Tokens," in "Proceedings of the 39th International Conference on Machine Learning (ICML)," Vol. 162 of *PMLR* 2022.

Brynjolfsson, Erik, "Information assets, technology and organization," *Management Science*, 1994, *40* (12), 1645–1662.

_ **and Andrew Ng**, "Big AI Can Centralize Decision-Making and Power—And That's a Problem," MILA-UNESCO Working Paper 2021.

_ **and Haim Mendelson**, "Information Systems and the Organization of Modern Enterprise," Working Paper 3560-93, MIT Sloan School of Management 1993.

_ **, Andrew McAfee, Michael Sorell, and Feng Zhu**, "Scale without Mass: Business Process Replication and Industry Dynamics," Harvard Business School Working Paper 07-016 2008.

_ **, Danielle Li, and Lindsey Raymond**, "Generative AI at Work," *The Quarterly Journal of Economics*, May 2025, *140* (2), 889–942. Open Access.

_ **, Wang Jin, and Xiupeng Wang**, "Information Technology, Firm Size, and Industrial Concentration," Working Paper 31065, National Bureau of Economic Research 2023.

Budish, Eric, Peter Cramton, and John Shim, "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market-Design Response," *The Quarterly Journal of Economics*, 2015, *130* (4), 1547–1621.

Buterin, Vitalik, Zoë Hitzig, and E Glen Weyl, "A flexible design for funding public goods," *Management Science*, 2019, *65* (11), 5171–5187.

Capitalist, Visual, "The Ten Largest U.S. Companies by Market Capitalization," `https://www.visualcapitalist.com/ranked-largest-us-companies-by-market-cap` 2025.

Carlsmith, Joseph, "How Much Computational Power Does It Take to Match the Human Brain?," Technical Report, Open Philanthropy 2020.

Carmichael, James, "Google Knows You Better Than You Know Yourself," *The Atlantic*, 2014.

Cemri, Mert, Melissa Z. Pan, Shuyi Yang, Lakshya A. Agrawal, Bhavya Chopra, Rishabh Tiwari, Kurt Keutzer, Aditya Parameswaran, Dan Klein, Kannan Ramchandran, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica, "Why Do Multi-Agent LLM Systems Fail?," 2025.

Choromanski, Krzysztof, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamás Sarlós, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Łukasz Kaiser, David Belanger, Lucy Colwell, and Adrian Weller, "Rethinking Attention with Performers," in "International Conference on Learning Representations (ICLR)" 2021.

Collective Intelligence Project, "Introducing the Collective Intelligence Project: Solving the Transformative Technology Trilemma through Governance R&D," White paper, Collective Intelligence Project 02 2023.

Csikszentmihalyi, Mihaly, *Flow: The Psychology of Optimal Experience*, New York: Harper & Row, 1990.

Cullen, Zoë, Danielle Li, and Shengwu Li, "Labor as Capital: AI and the Ownership of Expertise," Working paper (unpublished) September 2025.

Dao, Tri, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré, "FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness," in "Advances in Neural Information Processing Systems (NeurIPS)" 2022.

Decker, Ryan A. John C. Haltiwanger, Ron S. Jarmin, and Javier Miranda, "Changing Business Dynamism and Productivity," *American Economic Journal: Macroeconomics*, 2020, *12* (1), 297–326.

Delacroix, Sylvie and Neil D. Lawrence, "Bottom-up Data Trusts: Disturbing the 'One Size Fits All' Approach to Data Governance," *International Data Privacy Law*, 2019, *9* (4), 236–252.

Drago, Luke and Rudolf Laine, "The Intelligence Curse," Whitepaper, intelligence-curse.ai 2025.

Evans, Richard, Kelly O'Neill, and John et al. Jumper, "Protein Complex Prediction with AlphaFold-Multimer," *bioRxiv*, 2022.

Fedus, William, Barret Zoph, and Noam Shazeer, "Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity," *arXiv preprint arXiv:2101.03961*, 2021.

Fei-Fei, Li, Rob Fergus, and Pietro Perona, "One-Shot Learning of Object Categories," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2006, *28* (4), 594–611.

Firmani, Donatella, Marco Maiorino, Paolo Merialdo, and Elena Nieddu, "Towards Knowledge Discovery from the Vatican Secret Archives. In Codice Ratio – Episode 1: Machine Transcription of the Manuscripts," *arXiv preprint arXiv:1803.03200*, 2018.

Fishman, Charles, *The Wal-Mart Effect*, New York: Penguin, 2006.

Garicano, Luis, "Hierarchy and the Organization of Knowledge in Production," *Journal of Political Economy*, 2000, *108* (5), 874–904.

GDPR, EU, "Regulation (EU) 2016/679: General Data Protection Regulation," Official Journal of the European Union L 119 (4 May 2016) 2016.

Gilens, Martin and Benjamin I. Page, "Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens," *Perspectives on Politics*, 2014, *12* (3), 564–581.

Glaeser, Edward L., Giacomo A. M. Ponzetto, and Andrei Shleifer, "Why Does Democracy Need Education?," *Journal of Economic Growth*, 2007, *12* (2), 77–99.

Goldin, Claudia and Lawrence F. Katz, *The Race Between Education and Technology*, Cambridge, MA: Belknap Press / Harvard University Press, 2008.

Gorwa, Robert, Reuben Binns, and Christian Katzenbach, "Algorithmic Content Moderation: Technical and Political Challenges in the Automation of Platform Governance," *Big Data & Society*, 2020, *7* (1), 1–15.

Grossman, Gene M. and Elhanan Helpman, "Protection for Sale," *American Economic Review*, 1994, *84* (4), 833–850.

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

Gu, Albert and Tri Dao, "Mamba: Linear-Time Sequence Modeling with Selective State Spaces," *arXiv preprint arXiv:2312.00752*, 2023.

Harari, Yuval Noah, *Homo Deus: A Brief History of Tomorrow*, New York: Harper, 2017.

Hardjono, Thomas and Alex Pentland, "Data cooperatives: Towards a foundation for decentralized personal data management," *arXiv preprint arXiv:1905.08819*, 2019.

Hart, Oliver and John Moore, "Property Rights and the Nature of the Firm," *Journal of political economy*, 1990, *98* (6), 1119–1158.

Harvard Business School, "Mrs. Fields Cookies," Case Study #191-035 1990.

Hassan, Samer and Primavera De Filippi, "Decentralized Autonomous Organization," *Internet Policy Review*, 2021, *10* (2).

Hayek, Friedrich A., "The Use of Knowledge in Society," *American Economic Review*, 1945, *35* (4), 519–530.

Hendrycks, Dan, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt, "Measuring Massive Multitask Language Understanding," *arXiv preprint arXiv:2009.03300*, 2021.

Hestness, Joel, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md. Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou, "Deep Learning Scaling is Predictable, Empirically," *arXiv preprint arXiv:1712.00409*, 2017.

Hoffmann, Jordan, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre, "Training Compute-Optimal Large Language Models," *arXiv preprint arXiv:2203.15556*, 2022.

Hsu, Enshuo, Ioannis Malagaris, Yong-Fang Kuo, Rizwana Sultana, and Kirk Roberts, "Deep learning-based NLP data pipeline for EHR-scanned document information extraction," *JAMIA open*, 2022, *5* (2), ooac045.

Huang, Saffron and Sam Manning, "Here's How To Share AI's Future Wealth," *Noema Magazine*, 2025. Essay.

Institute for Local Self-Reliance, "Independent Business in the United States: Trends and Policy Solutions," 2019.

Jain, Piyush, Sean CP Coogan, Sriram Ganapathi Subramanian, Mark Crowley, Steve Taylor, and Mike D Flannigan, "A review of machine learning applications in wildfire science and management," *Environmental Reviews*, 2020, *28* (4), 478–505.

Jain, Shrey, Zoë Hitzig, and Pamela Mishkin, "Contextual Confidence and Generative AI," in "2025 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)" IEEE 2025, pp. 281–301.

Jameson, Fredric, "Walmart as Utopia," Verso Books blog essay 2016.

Jensen, Michael C. and William H. Meckling, "Specific and General Knowledge and Organizational Structure," in Lars Werin and Hans Wijkander, eds., *Contract Economics*, Oxford: Blackwell Publishers, 1992, pp. 251–274.

Jin, Qiao, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu, "PubMedQA: A Dataset for Biomedical Research Question Answering," in "Proceedings of EMNLP-IJCNLP" 2019, pp. 2567–2577.

Johnson, Jeff, Matthijs Douze, and Hervé Jégou, "Billion-scale Similarity Search with GPUs," *arXiv preprint arXiv:1702.08734*, 2017.

Joshi, Mandar, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer, "TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension," in "Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)" 2017, pp. 1601–1611.

Jumper, John, Richard Evans, and Alexander et al. Pritzel, "Highly Accurate Protein Structure Prediction with AlphaFold," *Nature*, 2021, *596*, 583–589.

Kandpal, Niyati, Ari Holtzman, Hao Zhou, Douwe Kiela, and Elizabeth Clark, "Large Language Models Struggle to Learn Long-Tail Knowledge," *arXiv preprint arXiv:2307.XXXX*, 2023.

Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei, "Scaling Laws for Neural Language Models," *arXiv preprint arXiv:2001.08361*, 2020.

Kwiatkowski, Tom, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov, "Natural Questions: A Benchmark for Question Answering Research," *Transactions of the Association for Computational Linguistics*, 2019, *7*, 452–466.

Kwon, Spencer Y., Yueran Ma, and Kaspar Zimmermann, "Technology, Market Power, and Firm Investment," *Journal of Financial Economics*, 2022, *146* (3), 777–799.

Kwon, Woosuk, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica, "Efficient Memory Management for Large Language Model Serving with PagedAttention," in "Proceedings of the 29th ACM Symposium on Operating Systems Principles (SOSP '23)" Association for Computing Machinery New York, NY, USA 2023.

Lake, Brenden M., Ruslan Salakhutdinov, and Joshua B. Tenenbaum, "Human-level Concept Learning through Probabilistic Program Induction," *Science*, 2015, *350* (6266), 1332–1338.

Lange, Oskar, "On the Economic Theory of Socialism: Part One," *The Review of Economic Studies*, October 1936, *4* (1), 53–71.

Langlotz, Curtis P., "Will Artificial Intelligence Replace Radiologists?," *Radiology* Blog Post 2019. Accessed July 2025.

LeCun, Yann, Léon Bottou, Yoshua Bengio, and Patrick Haffner, "Gradient-Based Learning Applied to Document Recognition," *Proceedings of the IEEE*, 1998, *86* (11), 2278–2324.

Lee, Hau L., "Walmart's Supply Chain: The Invention of Retail Logistics," *Harvard Business Review*, January 2006.

Lerner, Abba P., "Theory and Practice in Socialist Economics," *The Review of Economic Studies*, October 1938, *6* (1), 71–75.

Leviathan, Yair, Matan Kalman, and Yossi Matias, "Fast Inference from Transformers via Speculative Decoding," *arXiv preprint arXiv:2211.17192*, 2023.

Liu, Nelson F., Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang, "Lost in the Middle: How Language Models Use Long Context," *Transactions of the Association for Computational Linguistics*, 2024, *12.*

Loecker, Jan De, Jan Eeckhout, and Gabriel Unger, "The Rise of Market Power and the Macroeconomic Implications," *Quarterly Journal of Economics*, 2020, *135* (2), 561–644.

Lubna, Naveed Mufti, and Syed Afaq Ali Shah, "Automatic number plate Recognition: A detailed survey of relevant algorithms," *Sensors*, 2021, *21* (9), 3028.

Malone, Thomas W and Stephen A Smith, "Modeling the performance of organizational structures," *Operations Research*, 1988, *36* (3), 421–436.

Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Yolanda Gil et al., "Artificial Intelligence Index Report 2025," 2025. arXiv:2504.07139.

Mercor, "Mercor Opportunities Board," `https://work.mercor.com/explore` July 2025. Sample listings show expert rates—for example, $270 / hour for dermatologists and $30–60 / hour for plant specialists.

Morgan Stanley Research, "Small Caps Have Been a Big Story After Recessions," `https://www.msci.com/research-and-insights/blog-post/small-caps-have-been-a-big-story-after-recessions` 2023.

Morozov, Evgeny, "Digital Socialism? The Calculation Debate in the Age of Big Data," *New Left Review*, 2019, *116/117*, 33–67.

Myerson, Roger B, "Graphs and cooperation in games," *Mathematics of operations research*, 1977, *2* (3), 225–229.

NASA Jet Propulsion Laboratory, "Mars Communication Disruption and Delay Report," `https://www.lpi.usra.edu/lunar/strategies/resources/M2M-ACR2023_MarsCommunicationDisruptionDelay.pdf` 2023.

Nguyen, Xuan-Phi, Shrey Pandit, Revanth Gangi Reddy, Austin Xu, Silvio Savarese, Caiming Xiong, and Shafiq Joty, "SFR-DeepResearch: Towards Effective Reinforcement Learning for Autonomously Reasoning Single Agents," *arXiv preprint arXiv:2509.06283*, 2025.

OECD, *Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave*, Paris: OECD Publishing, 2020.

O'Keefe, Cullen, Peter Cihon, Ben Garfinkel, Carrick Flynn, Jade Leung, and Allan Dafoe, "The windfall clause: Distributing the benefits of AI for the common good," in "Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society" 2020, pp. 327–331.

Opam, Kwame, "Amazon plans to ship your packages before you even buy them," *The Verge.*

OpenAI, "GPT-4o System Card," Technical Report, OpenAI August 2024.

_ , "Introducing GPT-4.1 in the API," `https://openai.com/index/gpt-4-1/` April 2025. Accessed: 2025-09-29.

Ott, Henry W., *Electromagnetic Compatibility Engineering*, Hoboken, NJ: John Wiley & Sons, 2009.

Pava, Juan N., Caroline Meinhardt, Haifa Badi Uz Zaman, Toni Friedman, Sang T. Truong, Daniel Zhang, Elena Cryst, Vukosi Marivate, and Sanmi Koyejo, "Mind the (Language) Gap: Mapping the Challenges of LLM Development in Low-Resource Language Contexts," White Paper, Stanford Institute for Human-Centered Artificial Intelligence April 2025.

Peters, Jay, "Google Translate is Getting Support for More Than 110 New Languages," June 2024. The Verge report explaining that Google Translate's total rises to 243 languages.

Phillips, Leigh and Michal Rozworski, *The People's Republic of Walmart: How the World's Biggest Corporations Are Laying the Foundation for Socialism*, Verso Books, 2019.

Pichai, Sundar and Demis Hassabis, "Our next-generation model: Gemini 1.5," `https://blog.google/technology/ai/google-gemini-next-generation-model-february-2024/` February 2024. Accessed: 2025-09-29.

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

Posner, Eric and Eric Weyl, *Radical markets: Uprooting capitalism and democracy for a just society*, Princeton University Press, 2018.

Pozzolo, Andrea Dal, Giacomo Boracchi, Olivier Caelen, Cesare Alippi, and Gianluca Bontempi, "Credit Card Fraud Detection: A Realistic Modeling and a Novel Learning Strategy," *IEEE Transactions on Neural Networks and Learning Systems*, 2018, *29* (8), 3784–3797.

Radner, Roy, "The Organization of Decentralized Information Processing," *Econometrica*, 1993, *61* (5), 1109–1146.

_ **and Timothy Van Zandt**, "Information Processing in Firms and Returns to Scale," *Annals of Economics and Statistics*, 1992, (25–26), 265–298.

Richman, Tom, "How Mrs. Fields Cookies Taps Computers to Bake the Perfect Cookie," *Inc. Magazine*, September 1987. Cover story quoting Debbi Fields.

Schick, Timo, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom, "Toolformer: Language Models Can Teach Themselves to Use Tools," *arXiv preprint arXiv:2302.04761*, 2023.

Schroff, Florian, Dmitry Kalenichenko, and James Philbin, "FaceNet: A Unified Embedding for Face Recognition and Clustering," *arXiv preprint arXiv:1503.03832*, 2015. *Serial Acquisitions and Industry Roll-ups: Background Note*

Serial Acquisitions and Industry Roll-ups: Background Note*, Technical Report DAF/COMP(2023)13, OECD Competition Committee November 2023.*

Sevilla, Jaime and Edu Roldán*, "Training Compute of Frontier AI Models Grows by 4–5x per Year," Epoch AI Blog May 2024.*

Shepardson, David*, "Uber, Waymo launch autonomous ride-hailing service in Atlanta," Reuters, June 2025. Accessed 9 July 2025.*

Simon, Herbert A.*, "A Behavioral Model of Rational Choice," The Quarterly Journal of Economics, 1955, 69 (1), 99–118.*

Smith, Dominic A. and Sergio Ocampo*, "The Evolution of US Retail Concentration," American Economic Journal: Macroeconomics, 2025, 17 (1), 71–101.*

Spiegel, Joel R., Michael T. McKenna, Girish S. Lakshman, and Paul G. Nordstrom*, "Method and System for Anticipatory Package Shipping."*

Stanford HAI*, "AI Index Report 2025," Technical Report, Stanford University, Institute for Human-Centered Artificial Intelligence (HAI), Stanford, CA 2025. AI Index Steering Committee.*

Steiner, Christopher*, "Wall Street's Speed War,"* Forbes*, Sept. 27, 2010.* *https://www.forbes.com/forbes/2010/0927/outfront-netscape-jim-barksdale-daniel-spivey-wall-html 2010.*

Thompson, Nicholas*, "When Tech Knows You Better Than You Know Yourself," WIRED.*

von Mises, Ludwig*, "Die Wirtschaftsrechnung im sozialistischen Gemeinwesen," Archiv für Sozialwissenschaft und Sozialpolitik, 1920, 47, 86–121.*

_ *, "Economic Calculation in the Socialist Commonwealth," in Friedrich A. Hayek, ed., Collectivist Economic Planning: Critical Studies on the Possibilities of Socialism, London: George Routledge & Sons, 1935, pp. 87–130.*

Wood, Lina*, "How Europe's SMEs Proved Their Resilience During COVID-19,"* World Economic Forum*, Jan. 2023.* *https://www.weforum.org/stories/2023/01/europe-smes-business-resilience-covid-19/ 2023.*

Yao, Shunyu, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao*, "ReAct: Synergizing Reasoning and Acting in Language Models," in "International Conference on Learning Representations (ICLR)" 2023.*

Zuboff, Shoshana*, The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power, New York: PublicAffairs, 2019.*