AI's Use of Knowledge in Society

Erik Brynjolfsson*
Zoë Hitzig†

September 2025

Prepared for NBER Volume on the Economics of Transformative AI.
Preliminary and incomplete—comments welcome.

Abstract

Transformative AI invites a reappraisal of Hayek's "knowledge problem" and its implications for how decision rights are allocated within firms and society. We develop a simple property rights framework in which AI shifts the optimal locus of control through two channels: (i) by codifying and transferring once-tacit, previously inalienable local knowledge, and (ii) by expanding the capacity of central agents to aggregate, interpret, and act on interdependent data. These forces erode the informational advantage of on-the-spot decision makers, making centralized control more feasible and, in some settings, more efficient—especially where complementarities across information assets are strong. The framework yields several predictions: larger average firm size; greater industry concentration; and reduced local managerial autonomy in IT- and AI-intensive sectors. We review early evidence consistent with these patterns. We also discuss conditions that can still favor decentralization. Beyond the efficiency benefits of centralization, we highlight distributional and political-economy consequences: centralization can weaken workers' and communities' bargaining power, dampen incentives to invest in human capital, and amplify elite influence over policy and public discourse.

*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

F. A. Hayek's 1945 essay "The Use of Knowledge in Society" opens by criticizing a plan – favored by some of his contemporaries – for constructing a "rational economic order" from the top down (Hayek, 1945). He was responding mostly 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 shadow 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.

To some observers in Hayek's time, Lange and Lerner's writings seemed vindicated by the wartime economy. In many countries, emergency agencies had replaced market coordination with direct command. In the U.S., for example, the War Production Board set price ceilings, rationed raw materials, and dictated output schedules. And the results were impressive: between 1940 to 1944, total output of manufactured goods increased by 300 percent – a surge orchestrated almost entirely from Washington D.C. (Milward, 1979).

But to Hayek, the flaws of central planning had never been more evident. His criticism of his contemporaries was epistemic: it was 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." There are cognitive and physical limits on which forms of information can travel, and how, and how quickly. Washington may have coordinated a great deal of wartime production but it did so through communication channels that operated with weeks- or months-long delays. Much of what people know – all that is relevant to choosing the best economic outcomes – is local or specialized knowledge, tailored to the instant, or otherwise tacit and difficult to articulate. Any analysis that imagines otherwise, Hayek warned, simply "assumes the problem away."

The essay, along with Hayek's other writings of the time (Hayek, 1937, 1940), became a decisive contribution to the Socialist Calculation Debate. Hayek's epistemic objection was somewhat new even relative to von Mises' arguments – even if a planning board could discover prices (as Lange and Lerner contended), he argued, it would not be able to incorporate dispersed knowledge as well as a decentralized price system can. In other words, Hayek highlighted the impossibility of centralizing the information that prices embody, and argued that a decentralized market, where decisions are made locally by those who possess only partial knowledge, was the best way to a "rational economic order." In his words, "If

we can agree that the economic problem of society is mainly one of rapid adaptation to changes... the ultimate decisions must be left to the people who are familiar with these circumstances."

And the second half of the twentieth century seemed to endorse Hayek's view. The collapse of the Soviet-led planning bloc in 1989 and the subsequent wave of liberalization across the developing world turned Hayek's epistemic critique into near-consensus (Kornai, 1992; Stiglitz, 1996). As political scientist Francis Fukuyama remarked at the time, the moment marked "an unabashed victory of economic and political liberalism" (Fukuyama, 1992). By the early 2000s, economic historian Richard Nelson would write that "the close of the twentieth century saw a virtual canonization of market organization as the best, indeed the only effective way to structure an economic system" (Nelson, 2005).

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 arguments and thus their implications for the organization of economic decision-making. Hayek's arguments rest on the idea that it is impossible to gather, codify and transfer "knowledge of the particular circumstances of time and place" in real-time. This impossibility may have seemed to be a technological fact in his day. The War Production Board, for example, relied on punch-card tabulators and paper forms – hundreds of questionnaires mailed to thousands of firms – to learn what inventory or machine-time was available (Novick and Steiner, 1947). Against these technological realities, Hayek's claim – that the person on-the-spot always possesses knowledge a distant planner cannot – appeared self-evident.

Early information technologies slowly began to chip away at Hayek's central claim. IBM's MAPICS (Manufacturing, Accounting and Production Information Control System), introduced in 1977, allowed plant managers to upload data each night so corporate planners could recompute material requirements centrally. Enterprise software of the 1990s – like SAP R/3 – allowed headquarters to see production, purchasing and sales data in near-real time across entire multinational firms. Walmart's "Retail Link" gathered daily checkout data from every store at their headquarters in Bentonville. The result of the introduction of these technologies was a gradual movement of production and allocation decisions away from local managers and toward headquarters (Brynjolfsson and Mendelson, 1997; Acemoglu et al., 2007; Bloom et al., 2009).

AI further erodes Hayek's epistemic claims. It's not just that computerized databases are getting bigger and cheaper – continuing the trends that started in the 80s and 90s – but that machine learning initiates an entirely new paradigm. By codifying tacit, specialized expertise and processing vast, interdependent data streams at scale, AI turns once-local knowledge into centrally usable information, strengthening the possibilities for – and benefits of – centralizing decision rights. Consider, for example, Bosch's cloud-based AI Analytics Platform which inspects visual data from more than 2,000 production lines across roughly 50 plants, flagging anomalies from a single hub and guiding line-level fixes centrally. Or John Deere's "See & Spray" which streams camera feeds to a central command center, where aggregated weed maps generate updated, targeted spray prescriptions that sync back to connected machines in the field. Booming investment alongside rapid technical advances in AI are expanding the frontier of machine-codified and machine-processed knowledge rapidly, with little sign the momentum will slow (Maslej et al., 2025).

In this paper, we investigate the implications of the rise of AI for the organization of economic production, starting from the observation that many claims about the dispersed nature of knowledge – which appeared unquestionable in Hayek's day – are now in question. We offer a simple conceptual framework for understanding transformative AI's centralizing

tendencies, and early evidence of both the causes and consequences of this centralization (section 2). The analysis is based in a straightforward application of the property rights theory of the firm for a world (Grossman and Hart, 1986; Hart and Moore, 1990) combined with the changing alienabilty of information assets (Brynjolfsson, 1994). We model the rise of AI as affecting contracting through two channels: (i) which information assets are alienable, and (ii) increases in information processing capacity. We offer facts about recent developments in AI that justify our modeling choices.

Then, we discuss the effects of each channel in turn. First, when formerly inalienable information assets become alienable, the ownership of complementary assets shifts from local units that once held inalienable knowledge to whichever party makes the crucial investment, which in many cases may be the central hub (section 3). In other words, AI undermines Hayek's epistemic rationale for keeping decisions local. Second, when limits on information processing ease, centralizing control can more effectively harness complementarities between information assets (section 4). After considering forces that push in the opposite direction (section 5), we document some early evidence of increased centralization and concentration through these two channels (section 6). Finally, we discuss how the concentration of economic power can lead to the concentration of political power, and, without making any policy recommendations, we highlight how these dynamics could test the resilience of democratic institutions (section 7).

We are not the first to point out AI's potential for greater centralization of decision-making, but relative to the magnitude of the potential effects, there have been few economic analyses. As we discuss in more detail in section 6, there is an 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., 2023a), but these analyses have not yet incorporated the step-change that could result from transformative AI. Several writers have framed the observation that machine learning may centralize economic power in political rather than strictly economic terms.[1] Others have offered casual and speculative explorations of how AI enables centralization.[2] A few writings by economists share the spirit of this chapter (Brynjolfsson and Ng, 2021; Agrawal et al., 2022; Acemoglu, 2023).

2 Framework: Incomplete Contracts with Information Assets

We begin with a simple two party incomplete contracts framework (Grossman and Hart, 1986; Hart and Moore, 1990) in which some assets are information assets, following Brynjolfsson (1994). In this section we study the baseline case in which information assets are inalienable. Then, in section 3 and section 4, we model AI as affecting two aspects of the contracting environment: first, the alienability of information assets (i.e. whether an agent's knowledge or know-how can be codified and transferred) and second, the benefits and costs of information processing power.

[1] Writers on the contemporary left argue – 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. Frederic Jameson makes this case in his 2016 essay "Walmart as Utopia" (Jameson, 2016), and Leigh Phillips and Michal Rozworski expand on it in their 2019 book *The People's Republic of Walmart* (Phillips and Rozworski, 2019). A similar line of thought appears in Evgeny Morozov's essay "Digital Socialism" (Morozov, 2019).

[2] See, for example, the provocatively titled book *Fully Automated Luxury Communism* (Bastani, 2019), and a recent influential whitepaper which quips "Note... that money will continue being a thing, at least unless we have one single AI system doing all economic planning" (Drago and Laine, 2025).

We call the two parties the local "entrepreneur," E and the central "headquarters," H. There is a tangible asset a_F which can be owned by either agent, and an information asset a_E which is – in the baseline case – inalienable and in the mind of the entrepreneur. For example, H could be the central office of a cafe franchise and E could be a local manager of a cafe – the tangible asset a_F is the physical shop and ovens and a_E is the local manager's detailed understanding of neighborhood demand. This set up could also represent a field service dispatcher (H) and a technician (E) – the tangible asset a_F is the technician's van stocked with tools, and the information asset a_E is the technician's expertise in diagnosing and fixing unexpected faults on site.

We assume that contracts are incomplete – i.e. there are contingencies that the parties cannot foresee and investments are non-contractible. The timing is as follows. At stage 0, the parties choose (cooperatively) who owns the tangible asset a_F, which we represent through ownership parameter

$$\rho \in \{E, H\}.$$

At stage 1, the parties make non-contractible relationship-specific investments x_E and x_H and incur convex costs $c_E(x_E)$ and $c_H(x_H)$ in doing so. 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 is generated only if the tangible asset F and the information asset E are used together; they are strictly complementary.

More formally, surplus realized is given by $V(x_E, x_H, a_E, a_F)$ where V has the following properties: (i) it is strictly increasing in each party's investment, (ii) it exhibits diminishing returns to each party's investment, (iii) it exhibits weak complementarities in the two party's investments, and (iv) it is 0 whenever a_E or a_F is used without the other. Often we will suppress the dependence of V on the assets present, and write simply $V(x_E, x_H)$. After the output is realized, the two parties observe it and bargain over the surplus. We adopt Nash bargaining with bargaining weight α for E and $(1-\alpha)$ for H.[3] The party that owns a_F has a residual right of control – i.e. they can decide whether the asset is used in joint production if negotiations break down.

This residual right of control affects outside options. Because a_F is only valuable alongside the information asset a_E, when the entrepreneur does not own a_F, the outside options for both E and H are 0. In the cafe example, this means that a franchiser-owned outlet leaves the local manager with no viable independent fallback if a contract breaks down – she cannot leverage her neighborhood know-how to open her own shop. In the field-service example, a technician without title to the van and tools cannot continue working as an independent contractor with relationships she has cultivated. But when the entrepreneur does own a_F, their outside option is some increasing function of their investment, i.e. her outside option is $g_E(x_E)$ with $g_E(x_E) \geq 0$ for all x_E and $g'_E(x) > 0$. In the cafe case, owning the store and ovens means each increment of effort – say, mastering a new seasonal recipe or building relations with local suppliers – directly raises the value the manager could capture by going solo; similarly, a technician who holds title to her van and diagnostic kit boosts her alternative earnings by investing in client rapport.

Under these circumstances, when the information asset is inalienable, the standard Grossman-Hart-Moore logic suggests that it is better to co-locate decision rights and information – that is, the entrepreneur should own a_F. To see this, consider first the case where the entrepreneur does not own a_F, i.e. $\rho = H$. Denoting by V_i the marginal value

[3] Any efficient bargaining rule with ownership-based outside options delivers the same qualitative results.

of i's effort in joint production, the effort decisions of both parties (x_E, x_H) are determined by the follow conditions:

$$\alpha V_E(x_E, x_H) = c'_E(x_E) \tag{1}$$
$$(1-\alpha) V_H(x_E, x_H) = c'_H(x_H). \tag{2}$$

Next, consider the case where the entrepreneur owns a_F. The effort decisions of both parties (x_E, x_H) when $\rho = E$ are determined by:[4]

$$\alpha V_E(x_E, x_H) + (1-\alpha) g'_E(x_E) = c'_E(x_E) \tag{3}$$
$$(1-\alpha) V_H(x_E, x_H) = c'_H(x_H). \tag{4}$$

These two sets of first order conditions show that giving ownership of the tangible asset to the entrepreneur increases the incentives of the entrepreneur (compare (1) to (3)) while leaving the headquarters' incentives unchanged (compare (2) to (4)).

Fact 1. *If the information asset a_E is inalienable, then joint surplus is higher when the entrepreneur owns the tangible asset ($\rho = E$) than it is when the headquarters owns the asset ($\rho = H$).*[5]

This result formalizes Hayek's intuition that effective control should rest with the agent who possesses non-transferable, decision-relevant knowledge. When contracts are incomplete, incentives are provided by via the residual rights of control that come from asset ownership. Thus when a_E cannot be transferred, as Hayek posits is the case for local knowledge, giving the local party (E) 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. Knowing this is the ex-post outcome will raise the entrepreneur's ex-ante incentives for investment, and thus total surplus. Under franchiser ownership, that fallback is zero, so the manager has diminished incentives to invest in the local adaptations that could drive profits.

We will model transformative AI as affecting this simple setup through two channels. First, AI is codifying local, specialized, and tacit knowledge, making more "information assets" alienable. Second, AI lowers the cost and raises the benefits of central information processing – it relaxes the cognitive bottlenecks that once limited how much a hub could absorb, while aggregate-data learning makes every extra stream of data even more valuable. Each channel on its own pushes against Hayek's logic. Taken together, the two channels suggest potentially dramatic changes in asset ownership and control in the presence of transformative AI.

[4] Because the entrepreneur now owns the tangible asset, her outside-option payoff in bargaining is $g_E(x_E)$, whereas headquarters' outside option remains 0. With Nash bargaining weights $(\alpha, 1-\alpha)$ the ex-post pay-offs solve

$$\max_{u_E, u_H} (u_E - g_E)^{\alpha} u_H^{1-\alpha} \quad \text{s.t.} \quad u_E + u_H = V(x_E, x_H),$$

where u_E and u_H are ex-post payoffs of the entrepreneur and headquarters respectively. The first-order conditions give

$$u_E(x_E, x_H) = g_E(x_E) + \alpha\big[V(x_E, x_H) - g_E(x_E)\big], \qquad u_H(x_E, x_H) = (1-\alpha)\big[V(x_E, x_H) - g_E(x_E)\big].$$

At stage 1 each party chooses its investment to maximize continuation value minus cost: E solves $\max_{x_E} u_E - c_E(x_E)$ and H solves $\max_{x_H} u_H - c_H(x_H)$. Differentiating with respect investment decisions x_E and x_H yields (3) and (4).

[5] Formal statements and proofs of all facts are given in Appendix A.

3 AI Increases Codified Knowledge

In this section, we first discuss how AI is expanding the set of codifiable, and thus alienable, knowledge. Then we return to the GHM framework and analyze the implications of this expansion of codified knowledge for the location of economic decision-making.

We can disambiguate three trends:

1. The increased accessibility and digitization of physical records and other *explicit knowledge*,
2. The extraction and codification of *tacit knowledge*, and
3. The discovery of entirely new *machine-native knowledge* – never comprehended or comprehensible by any collection of human minds.

We discuss each of the trends in turn.

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[6] (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

[6] 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.

What are the implications for the allocation of decision rights? Let's return to the framework and analyze how the conclusions change when the information asset – once an inalienable asset held by the human entrepreneur – can now be codified and transferred to another party. The set of possible ownership structures is now

$$\rho^{\text{alienable}} \in \{(E,E),(E,H),(H,E),(H,H)\}$$

where each tuple (i,j) indicates that i owns a_E and j owns a_F.

We consider only arrangements in which both assets are owned by one party.[7] First, note that when both a_E and a_F are owned by the entrepreneur (i.e. $\rho = (E,E)$), E and H's choices of effort are still determined by (3) and (4) – this case is exactly the same as the $\rho = E$ case in the previous section. However, when the information asset can be transferred to H, we must consider a new ownership arrangement in which H owns both a_E and a_F. In this case, the investment decisions are determined by:

$$\alpha V_E(x_E, x_H) = c'_E(x_E) \tag{5}$$

$$(1-\alpha)V_H(x_E, x_H) + \alpha g'_H(x_H) = c'_H(x_H), \tag{6}$$

where $g_H(x_H)$ is H's outside option payoff when negotiations breakdown, with $g_H(x_H) \geq 0$ for all x_H and $g'_H > 0$.

[7]Since neither party has the ability to produce unless they own control both assets in the event that negotiations breakdown, the ownership arrangements (E,H) and (H,E) (in which the assets are split up) are dominated.

So, which ownership arrangement leads to higher joint surplus? Here, the ownership arrangement with higher joint surplus depends on which party's investment incentives matter more. The comparison hinges on how large the boost from incentives is relative to the share of output the party already internalizes. We have the following sufficient condition:[8]

Fact 2. *Suppose the information asset a_E is alienable. If for every feasible equilibrium investment level (x_E, x_H)*

$$(1-\alpha)V_H(x_E, x_H)g'_H(x_H) > \alpha V_E(x_E, x_H)g'_E(x_E), \tag{7}$$

then ownership of a_F and a_E by headquarters ($\rho = (H, H)$) delivers strictly higher surplus.

Put differently, inequality (7) says that both assets should be held by the side whose marginal investment, scaled by its bargaining weight, adds more to joint surplus. In many modern settings this may favor the central hub. Consider a franchised-café roll-up in which the private-equity sponsor installs a chain-wide AI demand-forecasting and procurement system. Each additional dollar of headquarters investment in model refinement and GPU capacity materially raises aggregate surplus—and, because those assets can be redeployed or sold, also enlarges the sponsor's outside option – so the left hand side of (7) is large. By contrast, a store manager's incremental know-how, once encoded in the central model, affects only her own outlet and has little standalone value, so the right hand side of (7) may be small.

This is quite a different conclusion from the conclusion of the baseline case in section 2. It suggests that decisions and information should be collocated, and sometimes that means moving the information instead of the decisions. This aligns with insights in Jensen and Meckling (1992) and Brynjolfsson and Mendelson (1993), and overturns 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." Contra Hayek, beneficial use can be made of the information asset a_E without leaving the decision to the entrepreneur.

4 AI Increases the Benefits of Centralization while Decreasing Costs

In the previous section, we saw that when information is alienable, ownership of the asset bundle goes to whichever party's marginal effort is more valuable. We noted some settings in which the central headquarters may have the more valuable investment, pushing in favor of centralization. In this section, we explore two further – and related – ways that developments in AI may push toward greater centralization: AI increases the benefits of centralization (through improved coordination and learning) and it eases the constraints on centralization through improved information processing capacity. We discuss each of these phenomena in turn.

4.1 AI Improves Coordination and Learning

Centralized control enables better coordination and learning. When decision-making authority is concentrated in a single decision-maker, "one individual takes responsibility for

[8] Note that (7) is a strong condition (it requires point-wise dominance). See subsection A.2 for a functional form example that leads us to a tight scalar necessary and sufficient condition.

the work of others, issuing instructions to them and monitoring their actions"—exemplifying what Mintzberg (1979) calls *direct supervision*. By contrast, splitting interdependent decisions among multiple principals invites frictions and misaligned incentives that erode performance (Bernheim and Whinston, 1986; Malone and Crowston, 1994). Modern machine-learning magnifies the gains from centralized coordination. A cloud optimizer can ingest real-time telemetry from hundreds of distributed agents, search an immense action space, and dispatch globally consistent commands every few milliseconds. Amazon's Kiva system, for instance, replans the paths of roughly 500 mobile robots in each fulfillment center, doubling packing throughput relative to human dispatchers (Wurman et al., 2008). In a very different setting, DeepMind's reinforcement-learning controller tunes hundreds of actuators in Google data centers, cutting cooling electricity by about 40 % (Evans and Gao, 2016).

Pooling data in one place also accelerates learning. Under the "cloud robotics" paradigm (Kuffner, 2010), every remote device becomes a sensor feeding a shared model. Waymo's vehicles stream their sensor traces to a common training pipeline so the fleet learns from billions of driven miles rather than any one car's experience (Kuffner, 2010; Pratt, 2015). Likewise, Netflix aggregates the viewing histories of more than 200 million subscribers; the resulting recommender models reliably outperform what any individual user could predict for herself (Gomez-Uribe and Hunt, 2015; Amatriain and Basilico, 2015). Together, these examples show how AI supercharges the longstanding coordination and learning advantages of centralization.

To represent these gains from coordination and learning, we return to our GHM framework and expand it. We continue, as in previous sections, to assume that there is a single central headquarters H. But we now suppose there is a set of N entrepreneurs, indexed by i, and each is endowed with an alienable information asset a_i. For each information asset a_i, there is a complementary tangible asset $a_{F(i)}$ such that a_i and $a_{F(i)}$ are only productive as a bundle – they must be used together or they produce zero output. As before, the entrepreneurs and the headquarters make non-contractible effort decisions $\mathbf{x} = (x_H, x_1, \ldots, x_N)$ at strictly convex costs $c_i(x_i)$.

Let A be the set of assets used in production. The production function $V(\mathbf{x}; A)$ depends on the number of bundles used in production and exhibits complementarities among information assets:

$$\frac{\partial^2 V}{\partial a_i \partial a_j} \geq 0, \quad i \neq j. \tag{8}$$

The tangible assets are not as complementary with each other as information assets are with each other, i.e. for a given i, j with $i \neq j$,

$$\frac{\partial^2 V}{\partial a_i \partial a_j} \geq \frac{\partial^2 V}{\partial a_{F(i)} \partial a_{F(j)}} \quad i \neq j. \tag{9}$$

Intuitively, owning two geographically distant vans yields little synergy beyond simple additivity, whereas owning two rich data streams from those vans can unlock non-linear value by improving forecasts across the whole fleet.

With many assets and many agents, the set of possible ownership arrangements expands (combinatorially). For simplicity, we will consider only ownership arrangements in which $K \leq N$ asset bundles are owned by the central hub, and we parametrize the ownership possibilities by K, i.e.

$$\rho^K \in \{0, 1, \ldots, N\}.$$

We call the extreme in which all N asset bundles are owned by the central hub "centralized ownership" (ρ^N) and the extreme in which all asset bundles are owned by the corresponding

entrepreneur "decentralized ownership" (ρ^0). We call intermediate ownership arrangements "partially centralized" (ρ^K).[9] We write $V(\mathbf{x}, K)$ to stress that the production function depends on the ownership arrangement.

When K assets are owned by the central hub, the value they get when negotiations breakdown $g_H(x_H, K)$ is increasing in the number of assets they control, i.e. $\frac{\partial g_H}{\partial K} \geq 0$ and, as before increasing in the investment x_H. The share of the surplus each agent receives when participating in production with K entrepreneurs and HQ is determined through a bargaining process—for ease of exposition, we assume the share they receive is α_i.[10] For every entrepreneur i who maintains ownership of her bundle, HQ and i enter a bilateral contract, in which case the entrepreneur receives a share α of the surplus (and HQ receives share $1 - \alpha$).

To more readily see the effects of the interdependencies between information assets, we make the complementarity between information assets explicit. We introduce an *interdependency parameter* $\lambda \geq 0$, defined such that:

$$\frac{\partial^2 V}{\partial a_i \partial a_j} = \lambda, \quad i \neq j. \tag{10}$$

This interdependency parameter captures the idea that greater learning can be done from diverse data when they are pooled, and greater coordination is possible when a single decision-maker holds interdependent information streams. Let $S(K, \lambda)$ be the total equilibrium surplus when HQ owns K bundles and the information complementarity parameter is λ.

Fact 3. *For every $K < N$, the incremental surplus from moving the $(K+1)^{st}$ bundle to headquarters,*

$$\Delta S(K, \lambda) = S(K+1, \lambda) - S(K, \lambda),$$

is weakly increasing in λ.

The remark formalizes the simple observation that when complementarities are stronger, the benefits to further centralization are greater. Consider Waymo's fleet of self-driving cars. The data from one car encountering a rare "edge case" – a complex construction zone or an erratic cyclist – is valuable. But its value is much greater when fed into a central model, instantly upgrading the entire fleet's capability and preventing thousands of potential future accidents. The value of centralizing one more car's data stream (increasing K) is amplified by the thousands already contributing. Similarly, for Netflix, the information asset is a user's viewing history. By itself, it offers modest value. But when pooled centrally, it enables the powerful collaborative filtering algorithms that discover that "people who watched *The Crown* and also liked the documentary *My Octopus Teacher* are unusually likely to enjoy the sci-fi series *Black Mirror*." The higher the complementarity (λ) the more powerful these

[9] It is worth noting that there are two further ownership arrangements that we do not analyze but are nonetheless easy to rule out. First, when HQ's effort is strictly productive (i.e. $\frac{\partial V}{\partial x_H} > 0$), ownership arrangements in which a coalition of entrepreneurs controls corresponding assets without HQ are dominated by the same coalition with HQ. Second, if we assume that asset bundles are more valuable in the hands of the entrepreneur they correspond to, we can rule out crossed-ownership arrangements, i.e. conditional on a single entrepreneur owning bundle i, it is better for entrepreneur i to own the bundle than any other j. There are other ownership arrangements that could be optimal but we do not study them here – for example, it could be that HQ owns K bundles and the remaining $N - K$ bundles could be owned by a coalition of the remaining entrepreneurs.

[10] As above, the precise share does not matter for our results.

discoveries become, and the stronger the economic rationale for pulling another user's data into the central system. As AI continues to advance, it effectively increases λ across more and more domains, strengthening the case for centralization.

4.2 AI Eases Limits on Information Processing Capacity

Despite the great benefits to centralization (via learning and coordination), these benefits have historically been challenging to harness. Indeed, 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).

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 10^15 FLOP/s, a threshold that modern supercomputers can now surpass, albeit at a high energy and monetary cost (Carlsmith, 2020). Meanwhile, the information processing capacity of AI has increasing exponentially. Beyond Moore's Law and its GPU-based counterparts, this is evident in the steep decline in the cost of computation, the well-documented "scaling laws" that link AI performance to computational resources (Kaplan et al., 2020; Sevilla et al., 2024; Ho et al., 2024), and the expansion of large language models' "working memory." In 2019, OpenAI's GPT-2 could process a context of 1,024 tokens; by 2024, models like Anthropic's Claude 3.5 could handle 200,000 tokens—the equivalent of a 500-page book—and by 2025, OpenAI's GPT4.1 could handle 1 million tokens.

The transformation in processing capacity is not merely about scale; it is about the fundamental architecture of knowledge acquisition and sharing. Herbert Simon observed that in human organizations, "it is usually important to specify where in the organization particular knowledge is stored... Depending on its actual locus, knowledge may or may not be available at the decision points where it would be relevant" (Simon, 1991). As Hinds and Pfeffer (2003) puts it, "organizations don't know what they know." This creates friction – information held by different people cannot be instantaneously accessed, introducing delays and misunderstandings incentive misalignment, and other communication frictions. As a result, organizations pay processing costs each time knowledge crosses a human boundary.

Contrast the dynamics of human organizations with a recent observation from Geoffrey Hinton, who noted that AI systems "can learn separately but share their knowledge instantly. So it's as if you had 10,000 people and whenever one person learnt something, everybody automatically knew it" (Hinton, 2023). This phenomenon of instantaneous knowledge pooling effectively creates a single actor with the combined learning ability of many, unbound by the communication delays that define human teams.

We formalize this in our GHM framework by introducing a limit on the information processing capacity. When a single actor handles multiple information assets, she incurs a cognitive cost. We model this cost with an upper bound, $\leq \bar{K} \leq N$, on the number of asset

bundles any one entity can effectively process. Let K be the number of bundles owned by the coalition operating the technology and set

$$\kappa := \min\{K, \bar{K}\}.$$

The production function and each owner's outside option depend on κ rather than on K itself:

$$V(\mathbf{x}, K) = \tilde{V}(\mathbf{x}, \kappa), \tag{11}$$

$$g_i(x_i, K) = \tilde{g}_i(x_i, \kappa). \tag{12}$$

Assets beyond $\bar{K}$ therefore add no value in production or in fallback payoffs.

Using (11)–(12), the first-order condition (FOC) for headquarters (HQ), which owns K bundles, is

$$\alpha_H \, \tilde{V}_H(\mathbf{x}, \kappa) + (1 - \alpha_H) \, \tilde{g}'_H\big(x_H, \kappa\big) \;=\; c'_H(x_H), \tag{13}$$

while an entrepreneur who *does* own her bundle satisfies

$$\alpha_i \, \tilde{V}_i(\mathbf{x}, \kappa) + (1 - \alpha_i) \, \tilde{g}'_i\big(x_i, 1\big) \;=\; c'_i(x_i), \tag{14}$$

and an entrepreneur who *does not* own her bundle satisfies

$$\alpha_i \, \tilde{V}_i(\mathbf{x}, \kappa) \;=\; c'_i(x_i). \tag{15}$$

Fact 4. *Let $1 < \bar{K} < N$ and consider an ownership pattern in which HQ already holds $\bar{K}$ bundles while every entrepreneur $j > \bar{K}$ owns her own bundle.*

1. *If HQ acquires one more bundle, joint surplus is weakly lower: HQ's FOC* (13) *is unchanged but entrepreneur j moves from* (14) *to* (15)*.*
2. *If instead the processing capacity rises from $\bar{K}$ to $\bar{K}+1$, giving HQ that extra bundle can now* increase *surplus, because the marginal asset enters both production and fallback payoffs.*

Hence a larger information-processing capacity enlarges the span of asset bundles that it may be efficient to centralize.

The bound $\bar{K}$ limits the returns to centralized ownership even when complementarities and incentive considerations would otherwise push for full centralization; expanding processing capacity relaxes this constraint and makes further centralization attractive.

5 Countervailing Forces

While we've laid out some reasons that powerful AI would lead to greater centralization of decision-making, there are three types of counterarguments that we should 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 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, preventing centralization. Third, even if we assume that AI is capable of all decisions and there are no legal or other reasons that humans need to be involved, there may still be a decentralization of decision-making among different AI entities.

Let's consider each of these cases in turn.

5.1 Limits to codifiability or machine information processing

As Hayek notes(Hayek, 1945), not all types of information are codifiable. Suppose that continues to be the case indefinitely. For instance, it has been argued that humans have an advantage in emotional intelligence. Detecting, even subconsciously, the subtle expressions on a customer's face or in their tone of voice can give a savvy salesperson an important advantage. If this information cannot be codified, it cannot be transmitted to a central decision-maker. In some cases, it may be advantageous to delegate decisions that require information that only humans can perceive to the human on the spot. Thus limitations on codifiability can lead to limitations on centralization.

Similarly, AI may fail for certain classes of decisions where humans do better. Consider ranking decisions from most common to least common. For instance, a self-driving car might frequently need to recognize and respond to a stop sign, while it might rarely, if ever, observe a black swan crossing the road. The distribution of cases typically reflects a power law, a few common cases, and a long tail with many potential situations for which there is no prior data. As a result, it can be much easier to gather data, the lifeblood of machine learning, for the common cases that make up the vast bulk of observations than for the rare cases for which there is little or no prior data.

In fact, Kandpal et al. (2023) have found that large language models struggle to learn long-tail knowledge – cases which only appear rarely in the training data. Agarwal et al. (2025) find that humans have a comparative advantage over AI in the long tail. While we humans are not great at solving problems that are rarely encountered, machine learning systems are usually even worse. A relative strength is that humans can generalize and extrapolate from very few examples (Fei-Fei et al., 2006; Lake et al., 2015). A toddler might recognize elephants after only being shown a few pictures, while current machine learning systems require much larger numbers of examples. Similarly, Langlotz (2019) argues that radiologists will not be replaced anytime soon by AI because they "know the long tail" of rare cases.

Similarly, humans currently have advantage in other areas, including types of perception and dexterity. As of this writing, competent humans know how to fold laundry, throw a curve ball, or button a shirt faster and more accurately than machines. This could require delegating those tasks, and the often-mundane decisionmaking associated with them, to humans rather than complete centralization.

That said, all these examples are fighting the premise that we were assigned for this conference. They assume that machines will not be able to match humans in at least some important capabilities. Over time, AI systems may obtain data for more cases, and more importantly, may learn to generalize from fewer cases, marching down the distribution and leaving human supremacy further and further out on the tail. One can contrast the performance of self-driving cars in 2012 to 2025, or the benchmarks like the MMLU in 2021 vs. 2024 to see that progress in AI can take some time, even if it's fairly monotonic.

5.2 Legislative requirements for decentralization

Suppose that machines become technically capable of making any decision as well or better than humans. Might economic decisions still be decentralized?

One reason this could be the case is that we, as a society, pass legislation that mandates some degree of decentralization. Today, anti-trust laws are invoked to forbid concentration of economic power and they could continue to be invoked even if greater centralization is more economically efficient.

Furthermore, even without legislation, there could be some people might simply prefer to work with human decisionmakers and these preferences could drive decentralization. As citizens or consumers might prefer to have a humans to be at the helm for certain classes of decisions: a human judge when life or liberty are at stake, a human pilot for critical flight decisions, a human player on the field for sports, and a human coach, poet or therapist for inspiration, art or advice. This would, de facto, lead to at least some decentralization.

While the focus of this article has been on the economic optimality of alternative arrangements for ownership and control, these decisions are not necessarily made on the basis of economics alone. As noted by Brynjolfsson and McAfee (2015) when "people see their income from labor recede, their views on the ownership of capital and the distribution of its proceeds, as expressed through votes or revolts, will matter even more than they do now."

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

Finally, let's consider the case where all decisions are made by AI. Even in this case, there may be reasons to decentralize decision-making.

First, machines can communicate vast quantities of information at the speed of light, but that speed is not infinite and communication is not always reliable. Thus, a local decision-maker may be able to respond (slightly) more rapidly and more dependably than a centralized one if they are physically separated. This is likely to be especially important when there are other AIs making local decisions that require strategic responses, especially adversarial ones. For instance, drone warfare may require millisecond responses in situations where communications are fragmented. Accordingly, decision-making autonomy and decentralization can create tactical advantages. Algorithmic trading can also require rapid reactions. Spread Networks installed a special low-latency fiber optic network between Chicago and New York because determined that shaving 3 milliseconds of the transmission time was worth the $300 million cost of the project (Steiner, 2010). Communications between Earth and Mars, with round-trip lags of 4.3 minutes to 21 minutes, necessitate significant autonomy in decision-making for the Mars landers and rovers (NASA Jet Propulsion Laboratory, 2023). When extreme speeds are required, even the short distance across a computer chip can be important. The latency may be measured in picoseconds/inch but that can matter at high clock speeds (Ott, 2009).

Similarly, while machines could eventually have vastly greater information processing capacity than humans, that capacity is not infinite or costless. There are both economies and diseconomies of scale in data centers. While larger data centers may have proportionately lower costs for cooling, redundancy, and power, they may also suffer from diseconomies due to bottlenecks and complexity. Where diseconomies of scale dominate, there will be an economic incentive to decentralize information processing and decision-making, even for decisions that are entirely allocated to machines.

Even among machines, contracts will be still be incomplete simply because the complexity of world makes it impossible to specify all contingencies in advance. The combinatorial possibilities of even moderately complex situations rapidly surpass even the fastest-growing exponential increases in processing power. Distributed agents, including AI agents, may need incentives to make non-contractible investments and in the above situations, this can require giving them residual rights or control. Thus, each member of a distributed swarm of AI agents, might be given some degree of asset ownership to preserve their decision-making autonomy.

5.4 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 we might reasonably expect that these cases account for a smaller and smaller share of the economy. Thus transformative AI implies orders of magnitude more centralization of decisionmaking. In short, a world that is radically different from the current one.

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. 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).

It's plausible that increased turbulence could be an important economic condition for some time, not least of which because of the uncertain effects of AI itself. In that case, which might not necessarily see increased centralization, at least during the transition toward increasingly powerful AI.

6 Early Empirical Evidence of Centralization

While artificial intelligence will likely become far more powerful in the coming years, there is already clear evidence that economic decision-making is shifting toward the center. 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).

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 Main-Street 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 – rather than within-market mergers – 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).

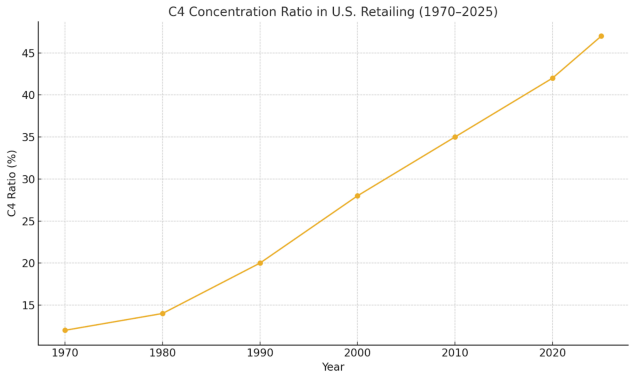


Figure 1: Growth of the U.S. retail C4 concentration ratio.
Source: Institute for Local Self-Reliance (2019).

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 (Autor et al., 2020). Stock-market data tell a similar story: the ten largest U.S. firms now account for roughly 38% of total market capitalization – double their share in 2010 (Capitalist, 2025).

A growing literature links these patterns to IT intensity. Brynjolfsson et al. (2023b) show that firms making heavier IT investments grow larger and more concentrated, consistent with earlier evidence in **?**. 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 these forces. 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.[11]

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. If these trends continue, we can expect more types of decisions to be made by machines rather than humans and economic activity to become more concentrated and centralized.

7 The Centralization of Economic Decisions and Political Power

The increasing centralization of economic decision-making and power enabled by AI may have important political consequences. Historically, the dispersion of economic power through competitive markets has been one of the foundations of liberal democracy, providing a counterweight to state authority and enabling broad-based participation in decision-making. The Cold War is a stark illustration: the relatively decentralized, competitive capitalist economies of the United States and its allies proved more resilient and adaptable than the centralized, command economies of the Soviet bloc. This suggests that economic decentralization can support political pluralism, while concentration of economic power risks reinforcing authoritarianism.

As Acemoglu and Johnson (2023) argues, technological advances often empower elites at the expense of broader societal interests unless countervailing institutions push toward more inclusive outcomes. If AI accelerates the centralization of economic decision-making—by concentrating data, computational resources, and strategic control in the hands of a few large actors—it could undermine the economic foundations that have historically supported democratic governance.

This concern is amplified by growing evidence that economic power translates into political influence. Gilens and Page (2014) finds that U.S. policy outcomes are strongly aligned with the preferences of economic elites and organized interests, but not with those of average citizens—implying that concentrated economic power can distort democratic responsiveness. More recently, Rau and Stokes (2024) documents that rising income inequality and economic concentration can erode democratic norms and increase support for authoritarian policies.

In light of these findings, the centralizing tendencies of AI could exacerbate political inequality, making it harder to sustain democratic checks and balances in the face of growing corporate and elite power. Just as decentralization of economic power contributed to the success of liberal democracies in the 20th century, renewed concentration of economic power in the 21st—driven by AI and digital platforms—could tilt the global battle of ideas in favor of more centralized, authoritarian models of governance.

The growing shift of decision-making to AI also risks suppressing incentives to invest in human capital, especially education. If students anticipate being outpaced or replaced by AI's superior problem-solving and information-processing capabilities, the perceived return on educational efforts falls. Early qualitative and emerging quantitative evidence provides

[11] This would be an effort to pass the Suleyman Test, and modern version of the Turing Test to create "artificial capable intelligence." Specifically, give an AI some capital, and a goal of making $ 1,000,000 autonomously, including "hierarchical planning, stitching multiple goals and subgoals and capabilities... to lobby, sell, manufacture, hire, plan—everything a company can do" (Suleyman, 2023).

some support for this. Livingstone (2023) recounts quitting her job teaching writing because "with the easy temptation of AI, many... of my students were no longer willing to push through discomfort." Similarly, a high-school English teacher recently reported that students "depend heavily on ... AI tools" in ways that lead to "apathy toward learning" (Maria, 2024). This growing apathy –combined with early research showing diminished critical thinking, effort, and motivation when AI is used (Holmes, 2023) – suggests a looming collapse in the value and accumulation of human capital.

This erosion of education has implications for democracy. Empirical research has found a link between education and democracy: human capital cultivates the civic knowledge, informed judgment, and critical engagement necessary for democratic citizenship (Dewey, 1916; Dahl, 1971). When education falters, democratic resilience deteriorates—inequality in political participation rises, public discourse is degraded, and authoritarian movements gain strength. In this light, declining motivation to learn—if triggered by over-dependence on AI—could reduce not only individual opportunity but also our collective capacity to sustain democratic societies. Without robust human capital, the incentives to think independently, evaluate policy, and hold power to account wane.

AI may also alter the dynamics of information control and public discourse. Providers of large language models increasingly serve as intermediaries in shaping the information environment, influencing news visibility, advertising markets, and even real-time political messaging (Zuboff, 2019; Gorwa et al., 2020). The design and deployment of these systems effectively grant their operators gatekeeping power over the flow of information, with potential implications for public opinion formation and electoral competition. This aligns with empirical evidence demonstrating that changes in the informational gatekeeping of social media platforms can shift electoral outcomes, as documented by Allcott and Gentzkow (2017) in their study of misinformation and voting behavior.

More broadly, the concentration of economic power facilitated by AI may strengthen firms' capacity to influence regulatory and legislative processes. The classical literature on special-interest politics predicts that greater concentration of cash flows and profits provides actors with additional resources for agenda setting, lobbying, and participation in rule-making (Grossman and Helpman, 1994)). In settings where economic rents become more concentrated in a small number of firms, these dynamics are likely to become more salient, further reinforcing the feedback loop between economic and political power.

8 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 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 favored local autonomy; second, by dramatically increasing information-processing capacity, AI enables greater integration and coordination of decisions across larger scales. Together, these channels diminish the informational advantages previously held by distributed human agents, thus making centralized control more feasible and, in some contexts, more efficient.

Our conceptual framework formalizes the intuition that, when critical knowledge assets are inalienable, it is optimal to colocate decision rights with the human agent who possesses

that knowledge. However, as codification advances and tacit or localized knowledge becomes transferable, ownership and control have the potential to shift toward the central hub, which can then better exploit complementarities across multiple information assets. Similarly, as AI systems expand the processing span of a single decision-making entity, the traditional argument for decentralization rooted in bounded rationality weakens.

These theoretical results 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 remove discretion from local managers in retail and other industries. This suggest that centralization is likely to intensify as AI capabilities advance

We also note that the centralizing tendencies of AI extend beyond efficiency considerations—to potentially reshape 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. At a societal level, reduced education and participation risk undermining civic engagement and democratic resilience, while concentrated economic resources and control over information flows may amplify the influence of elites over policy and public discourse.

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 maintain the foundations of democracy. Understanding and addressing these questions will 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 Simon Johnson, *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, New York: PublicAffairs, 2023.

_ , Philippe Aghion, Claire Lelarge, John Van Reenen, and Fabrizio Zilibotti, "Technology, Information, and the Decentralization of the Firm," *Quarterly Journal of Economics*, 2007, *122* (4), 1759–1799.

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.

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

Amatriain, Xavier and Justin Basilico, "Recommender Systems in Industry: A Netflix Case Study," in Francesco Ricci, Lior Rokach, and Bracha Shapira, eds., *Recommender Systems Handbook*, Boston, MA: Springer, 2015, pp. 385–419.

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

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.

Bernheim, B. Douglas and Michael D. Whinston, "Common Agency," *Econometrica*, 1986, *54* (4), 923–942.

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.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen, "The Organization of Firms across Countries," *Quarterly Journal of Economics*, 2009, *127* (4), 1663–1705.

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.

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

_ and Andrew McAfee, *The Second Machine Age*, New York: W. W. Norton, 2015.

_ 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.

_ and _ , "Information Systems and the Organization of Modern Enterprise," *Journal of Organizational Computing and Electronic Commerce*, 1997, *7* (1), 1–16.

_ , 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.

_ , _ , and Zi Wang, "Information Technology and the Growth of Large Firms," *Management Science*, 2023. Forthcoming.

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.

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.

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

Dahl, Robert A., *Polyarchy: Participation and Opposition*, New Haven: Yale University Press, 1971.

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.

Dewey, John, *Democracy and Education: An Introduction to the Philosophy of Education*, New York: Macmillan, 1916.

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

Evans, Richard and Jim Gao, "DeepMind AI Reduces Google Data Centre Cooling Bill by 40%," `https://deepmind.google/discover/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-by-40/` July 2016. DeepMind Blog.

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

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.

Fukuyama, Francis, *The End of History and the Last Man*, New York: Free Press, January 1992.

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

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.

Gomez-Uribe, Carlos A. and Neil Hunt, "The Netflix Recommender System: Algorithms, Business Value, and Innovation," *ACM Transactions on Management Information Systems*, 2015, *6* (4), 13:1–13:19.

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.

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.

Hayek, Friedrich A., "Economics and Knowledge," *Economica*, 1937, *4* (13), 33–54.

_ , "Socialist Calculation: The Competitive "Solution"," *Economica*, May 1940, *7* (26), 125–149.

_ , "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.

Hinds, Pamela J. and Jeffrey Pfeffer, "Why Organizations Don't 'Know What They Know': Cognitive and Motivational Factors Affecting the Transfer of Expertise," in Mark S. Ackerman, Volkmar Pipek, and Volker Wulf, eds., *Sharing Expertise: Beyond Knowledge Management*, MIT Press, 2003, pp. 3–26.

Hinton, Geoffrey, "Interview with BBC News on the Risks of Advanced AI," `https://www.bbc.com/news/world-us-canada-65452940` 2023.

Ho, Anson, Tamay Besiroglu, Pablo Villalobos, and et al., "Algorithmic Progress in Language Models," *arXiv preprint arXiv:2403.05812*, 2024.

Holmes, Wayne, "The Unintended Consequences of Artificial Intelligence and Education," Technical Report, Education International, Brussels October 2023.

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.

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.

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.

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, and et al., "Scaling Laws for Neural Language Models," *arXiv preprint arXiv:2001.08361*, 2020.

Kornai, János, *The Socialist System: The Political Economy of Communism*, Oxford: Oxford University Press, 1992.

Kuffner, James J., "Cloud-Enabled Robots," in "Workshop "What's Next?" at the IEEE-RAS International Conference on Humanoid Robots (HUMANOIDS 2010)" Nashville, TN 2010. Keynote talk.

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.

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.

Livingstone, Victoria, "I Quit Teaching Because of ChatGPT," *Time*, 2023. September 30. Available at `https://time.com/7026050/chatgpt-quit-teaching-ai-essay/`.

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 Kevin Crowston, "The Interdisciplinary Study of Coordination," *ACM Computing Surveys*, 1994, *26* (1), 87–119.

Maria, Hannah, "Students Depend Heavily on AI Tools," TikTok video 2024. Viral clip retrieved from SmartPicko, June 2025.

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.

Milward, Alan S., *War, Economy and Society, 1939–1945*, University of California Press, 1979.

Mintzberg, Henry, "An emerging strategy of" direct" research," *Administrative science quarterly*, 1979, *24* (4), 582–589.

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.

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

Nelson, Richard R., *Technology, Institutions, and Economic Growth*, Harvard University Press, 2005.

Novick, David and George A. Steiner, *Wartime Industrial Statistics: A Study of the War Production Board's Industrial Statistical Reporting Experience*, U.S. Government Printing Office, 1947.

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

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.

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

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.

Pratt, Gill A., "Is a Cambrian Explosion Coming for Robotics?," *Journal of Economic Perspectives*, 2015, *29* (3), 51–60.

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.

Rau, Eli G. and Susan C. Stokes, "Income Inequality and the Erosion of Democracy in the Twenty-First Century," *Proceedings of the National Academy of Sciences*, 2024, *122* (1), e2422543121.

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

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, Lennart Heim, Pablo Villalobos, and et al., *"Compute Trends Across Three Eras of Machine Learning,"* `https://epochai.org/blog/compute-trends` *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.*

_ , *"Organizations and Markets," Journal of Economic Perspectives, 1991, 5 (2), 25–44.*

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

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.*

Stiglitz, Joseph E., *Whither Socialism? Wicksell Lectures, Cambridge, MA: The MIT Press, January 1996.*

Suleyman, Mustafa, *"My New Turing Test Would See if AI Can Make $1 Million," MIT Technology Review, July 2023.*

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.*

Wurman, Peter R., Raffaello D'Andrea, and Mick Mountz, *"Coordinating Hundreds of Cooperative, Autonomous Vehicles in Warehouses," AI Magazine, 2008, 29 (1), 9–20.*

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

A Formal Statements of Facts

A.1 Remark 1

A.2 Remark 2

We give a functional form example to yield a stronger condition.

$$V(x_E, x_H) = x_E\, x_H, \qquad c_i(x_i) = \tfrac{1}{2}x_i^2, \qquad g_E(x_E) = \beta_E x_E, \quad g_H(x_H) = \beta_H x_H, \qquad \beta_E, \beta_H > 0.$$

Entrepreneur owns both assets ($\rho = (E, E)$)**.** Under Nash bargaining the continuation pay–offs are

$$u_E = g_E(x_E) + \alpha\big[V(x_E, x_H) - g_E(x_E)\big], \qquad u_H = (1-\alpha)\big[V(x_E, x_H) - g_E(x_E)\big].$$

Maximizing $u_i - c_i$ gives the first-order conditions

$$\alpha\, \partial_{x_E} V(x_E, x_H) + (1-\alpha) g_E'(x_E) = x_E, \qquad (1-\alpha)\, \partial_{x_H} V(x_E, x_H) = x_H.$$

Substituting the functional forms yields the equilibrium investments

$$x_E^* = \frac{(1-\alpha)\beta_E}{1-\alpha+\alpha^2}, \qquad x_H^* = \frac{(1-\alpha)^2\beta_E}{1-\alpha+\alpha^2}.$$

Headquarters owns both assets ($\rho = (H, H)$)**.** By symmetry,

$$x_E^{**} = \frac{\alpha^2\beta_H}{1-\alpha+\alpha^2}, \qquad x_H^{**} = \frac{\alpha\beta_H}{1-\alpha+\alpha^2}.$$

Total surplus.

$$W_{EE} = x_E^* x_H^* - \tfrac{1}{2}\big[(x_E^*)^2 + (x_H^*)^2\big], \qquad W_{HH} = x_E^{**} x_H^{**} - \tfrac{1}{2}\big[(x_E^{**})^2 + (x_H^{**})^2\big].$$

A little algebra shows

$$W_{HH} - W_{EE} = \frac{\alpha^2(1-\alpha)^2}{2\big(1-\alpha+\alpha^2\big)}\,(\beta_E - \beta_H)(\beta_E + \beta_H).$$

Hence headquarters ownership yields higher joint surplus iff $\beta_H < \beta_E$; entrepreneur ownership is superior when $\beta_H > \beta_E$. In this quadratic-cost, linear-outside-option benchmark the sign of $\beta_E - \beta_H$ fully determines the efficient allocation of the tangible and (now alienable) information assets.

A.3 Remark 3

A.4 Remark 4