

AI, Central Planning, and Hayek’s Knowledge Problem: Why “The Use of Knowledge in Society” Survives the Age of Artificial Intelligence

Heng-Fu Zou

April 2025

Abstract

This paper reassesses F. A. Hayek’s “The Use of Knowledge in Society” in light of the rise of artificial intelligence. AI substantially strengthens the case for certain forms of planning by expanding data collection, prediction, codification, monitoring, logistics optimization, and hierarchical coordination. Recent arguments by Acemoglu and by Brynjolfsson and Hitzig suggest that AI may reduce the informational advantage of local actors, enlarge the efficient scale of firms, and revive some forms of algorithmic or AI-assisted planning. The paper argues, however, that AI does not fundamentally invalidate Hayek’s knowledge problem. Hayek’s central claim was not merely that planners lack computing power, but that economically relevant knowledge is dispersed, tacit, contextual, incentive-laden, forward-looking, and often generated only through decentralized action. AI can process data, but data are not identical to knowledge; AI can predict, but prediction is not discovery; AI can simulate prices, but simulated prices are not market prices generated by property, exchange, competition, profit, and loss. The paper further argues that comprehensive AI planning raises serious problems of truthful revelation, metric gaming, surveillance, political power, and objective-function selection. AI may improve planning within firms, platforms, and governments, but such planning works best when embedded in a competitive, price-guided, polycentric market order. The conclusion is not anti-AI but anti-monopoly over social knowledge: AI should augment decentralized discovery rather than replace it with centralized control.

Keywords: Artificial intelligence; Hayek; knowledge problem; socialist calculation debate; central planning; market prices; tacit knowledge; dispersed knowledge; entrepreneurship; economic calculation; algorithmic governance; surveillance; firms; platforms; Coase; Mises; polycentric order.

JEL Classification: B25, B31, B53, D46, D80, D83, D85, L14, L16, L22, L23, O31, O33, P11, P16, P21, P51, P52.

1 Introduction

The extraordinary rise of artificial intelligence has reopened one of the oldest and most important debates in modern political economy: whether a complex economy can be rationally organized by a central intelligence. For almost a century, the standard liberal and Austrian answer to this question has been shaped by Ludwig von Mises's critique of socialist calculation and F. A. Hayek's theory of dispersed knowledge. Mises argued that without private property in the means of production, there could be no genuine market prices for capital goods and therefore no rational economic calculation (Mises, 1920, 1922, 1949). Hayek then deepened the argument by insisting that the central problem of economic order is not simply one of computation but one of knowledge: the relevant knowledge for coordinating production and consumption is not given to any single mind, planner, ministry, or statistical office (Hayek, 1937, 1945). It exists in fragmented, local, tacit, practical, and rapidly changing forms among millions and now billions of individuals. The real economic problem is therefore not merely how to solve an enormous allocation problem once all the data have been gathered; it is how to make use of knowledge that is never fully gathered in the first place.

Hayek's classic 1945 essay, "The Use of Knowledge in Society," formulated this point with unmatched clarity. He began by asking what problem society must solve when it attempts to construct a rational economic order. If all relevant information were already available to one mind, if preferences were given, and if the means of production were fully known, then allocation would be a problem of logic, mathematics, or engineering. But for Hayek this assumption already eliminates the real problem. The actual economic problem arises because "the data" are not given to anyone in their totality (Hayek, 1945). The knowledge that matters is dispersed among separate individuals, and much of it concerns what Hayek famously called "the particular circumstances of time and place" (Hayek, 1945). The knowledge of a local merchant, farmer, engineer, worker, consumer, shipper, entrepreneur, or inventor is not simply a small piece of an already completed social database. It is knowledge embedded in a situation, often known only through practice, judgment, and immediate action. The price system is valuable precisely because it allows people to adjust their plans to one another without requiring any one person to know the whole structure of social production (Hayek, 1945). Hayek's essay is therefore not merely a defense of prices as incentives; it is a theory of prices as a communication system, a discovery procedure, and a means of economizing on knowledge itself.

Artificial intelligence appears, at first sight, to place this argument under serious pressure. Hayek wrote in 1945, in a world of paper records, slow statistical reporting, primitive computing machines, and bureaucratic ministries unable to process the information already available to them. The economic world of the twenty-first century is fundamentally different. Digital platforms collect real-time data on purchases, searches, movements, inventories, reviews, logistics, payments, weather, employment, production, health, and communication. Machine learning systems detect patterns that human beings cannot see. Large language

models extract meaning from unstructured text, summarize documents, translate languages, classify behavior, write code, and assist in decision-making. Predictive algorithms forecast demand and shortages. Optimization systems reroute trucks, manage warehouses, schedule workers, price goods, detect fraud, and allocate attention. A planner armed with artificial intelligence would not resemble the Soviet bureaucrat imagined by Hayek’s generation. Such a planner would command cloud computing, sensor networks, data centers, autonomous software agents, predictive models, and perhaps real-time access to the behavior of whole populations.

For this reason, the AI challenge to Hayek must be taken seriously. It is intellectually inadequate simply to repeat that central planners cannot process enough information. AI greatly expands information-processing capacity. In some domains, it plainly gives central authorities, corporate headquarters, logistics networks, and digital platforms more knowledge than local decision-makers possess. A large retailer may use AI to predict demand across thousands of stores better than any individual store manager. A global shipping firm may reroute vessels and trucks more efficiently from a central operations center than through purely local judgment. A platform may infer consumer preferences from millions of transactions more accurately than a neighborhood seller can. A hospital system may use AI to forecast staffing needs, detect disease patterns, and allocate scarce equipment more effectively than isolated administrators. In these settings, AI does not abolish Hayek’s problem, but it does change the empirical boundary between what can be centralized and what must remain decentralized. This point is consistent with Coase’s older argument that firms arise because some forms of coordination can be performed more cheaply within hierarchy than through market contracting (Coase, 1937), and with Williamson’s later analysis of markets and hierarchies as alternative governance structures (Williamson, 1975, 1985).

The strongest recent formulation of this pro-AI challenge is Erik Brynjolfsson and Zoë Hitzig’s “AI’s Use of Knowledge in Society.” Their title deliberately echoes Hayek. They argue that artificial intelligence may codify forms of local or tacit knowledge that previously gave on-the-spot decision-makers an informational advantage. Where tacit knowledge can be extracted from behavior, text, images, voice, transactions, demonstrations, or digital traces, AI may make it transferable to headquarters, platforms, or centralized organizations. Brynjolfsson and Hitzig therefore argue that AI can shift the efficient locus of control, enlarge firm size, increase industry concentration, and reduce the need for local managerial autonomy in some sectors (Brynjolfsson and Hitzig, 2025). This is an important and sophisticated challenge because it does not crudely claim that Hayek was simply wrong. Rather, it says that the technological conditions under which Hayek’s argument operated have changed. If artificial intelligence can convert some tacit knowledge into usable code, models, and predictions, then the informational case for decentralization may weaken in some areas.

Daron Acemoglu has similarly asked whether AI-enabled centralized control could revive the possibility of effective planning. He recognizes that AI’s ability to absorb, organize, and interpret data on a massive scale may appear to strengthen the case for “AI socialism”

or some form of algorithmic economic planning (Acemoglu, 2023a, 2023b). If the older objection to planning was that bureaucrats could not collect and process enough information, then AI seems to alter the terms of the debate. Acemoglu nevertheless emphasizes that Hayek’s argument remains powerful because central planners still cannot adequately collect, interpret, and act upon the dispersed and often contradictory bits of knowledge held by individuals. They would need not only records of what people bought yesterday, but also knowledge of what people may want tomorrow, how they might use their talents creatively, and what new opportunities they may discover if left free to experiment (Acemoglu, 2023a). In this respect, the AI debate does not replace Hayek’s problem; it restates it under new technological conditions.

The purpose of this paper is to argue that AI revises Hayek’s argument but does not refute it. It is true that AI increases the feasibility of many forms of planning. It can improve firm organization, logistics, forecasting, public administration, scientific discovery, and even some forms of industrial policy. It can make the modern state more capable and the modern corporation more centralized. It can convert some previously tacit knowledge into data, models, or routines. It can reduce certain transaction costs and permit hierarchical organizations to coordinate activities at larger scale. For these reasons, a serious Hayekian political economy must acknowledge that the boundary between market coordination and organizational planning is historically variable. Coase’s theory of the firm already taught that firms exist because some coordination can be performed more cheaply within hierarchy than through market contracting (Coase, 1937). Williamson later developed this insight into a broader theory of transaction-cost economizing and governance (Williamson, 1975, 1985). AI changes the relative costs of market and hierarchy. In that sense, it may alter the optimal size of firms and the scope of internal planning.

But this concession is far from surrender. Hayek’s deepest argument was not that planners lacked sufficiently powerful calculating machines. His argument was that the central knowledge needed for economic coordination is not equivalent to objective, pre-existing, collectible data. Economic knowledge is dispersed, tacit, subjective, contextual, embodied, changing, and often generated only through the process of competition itself (Hayek, 1937, 1945, 1948). Artificial intelligence works best where data are abundant, feedback is rapid, tasks are repeated, objectives are clear, and patterns are stable enough to be learned from the past. Market society, however, is not merely a machine for allocating known goods among known uses. It is also a process for discovering new goods, new uses, new technologies, new preferences, new talents, new organizational forms, and new possibilities. This discovery process cannot be fully replaced by prediction because its results do not yet exist as data. This is why Hayek’s later formulation of competition as a “discovery procedure” remains central to the AI debate (Hayek, 1948; Hayek, 1968). It is also why Kirzner’s theory of entrepreneurial discovery remains indispensable: markets do not merely allocate known resources; they reveal unnoticed opportunities through alertness, rivalry, and profit-seeking action (Kirzner, 1973, 1985).

This distinction between prediction and discovery is central to the argument of the

paper. AI can forecast the future from the past, but entrepreneurial capitalism partly creates a future that cannot be inferred from the past. Before the personal computer, the smartphone, the search engine, cloud computing, mRNA vaccine platforms, generative AI, or countless other innovations, the relevant “demand” was not simply hidden in a database waiting to be discovered by a more powerful algorithm. It emerged through experimentation, persuasion, imagination, investment, competition, and use. Consumers often do not know what they want until entrepreneurs show them something new. Workers often do not know where their talents will be most productive until institutions permit mobility and experimentation. Firms often do not know which production method will dominate until rivals test alternatives. The market order is therefore not merely a computational device. It is an evolutionary and experimental order. This point links Hayek’s knowledge theory to Schumpeter’s theory of innovation and “creative destruction” (Schumpeter, 1911, 1942), Nelson and Winter’s evolutionary theory of economic change (Nelson and Winter, 1982), and Lavoie’s reinterpretation of the socialist calculation debate as a problem of rivalry, discovery, and process rather than mere computation (Lavoie, 1985a).

The same point applies to tacit knowledge. Michael Polanyi’s famous claim that “we can know more than we can tell” remains essential for understanding the limits of AI (Polanyi, 1966). AI can now infer some tacit patterns from behavior and can automate some skills without requiring explicit verbal articulation. This is a genuine technological breakthrough and a real challenge to older accounts of tacitness. David Autor’s discussion of Polanyi’s paradox is especially relevant here: tasks that were once hard to automate because people could not fully explain how they performed them may become automatable when machines learn from examples rather than from explicit rules (Autor, 2014). Brynjolfsson and Hitzig rightly note that modern AI can help overcome some versions of Polanyi’s paradox (Brynjolfsson and Hitzig, 2025). But it does not follow that all tacit knowledge is now codifiable. Much human knowledge remains inseparable from responsibility, judgment, trust, taste, context, timing, moral evaluation, and willingness to bear consequences. An AI system may infer a pattern, but the economic agent must still decide whether the pattern matters, whether it should be acted upon, and what risk to bear in acting upon it. Tacit knowledge is not only unarticulated information; it is often situated competence exercised under uncertainty (Polanyi, 1958, 1966; Hayek, 1945).

The paper therefore distinguishes three claims that are often confused. The first claim is that AI can improve planning. This is true. The second claim is that AI can improve central planning relative to older bureaucratic planning. This is also true. The third claim is that AI can make comprehensive central planning superior to decentralized market coordination. This is not established and, I argue, remains false. The reason is that comprehensive planning requires more than information processing. It requires genuine prices, incentive-compatible revelation of knowledge, entrepreneurial discovery, pluralism of ends, and political safeguards against concentrated power. AI can assist with calculation, but it cannot by itself create the institutional conditions under which calculation becomes meaningful. The older socialist calculation debate remains relevant precisely because it concerned not only mathematical

feasibility but also the institutional preconditions of meaningful economic calculation (Mises, 1920; Hayek, 1935, 1940, 1945; Lange, 1936, 1937; Lavoie, 1985a).

Here the Misesian argument remains indispensable. Even if a planning agency possessed vast data, it would still face the problem of economic calculation in the absence of genuine markets for capital goods. Market prices are not merely numbers that can be simulated by an algorithm. They emerge from exchange under private property, scarcity, competition, profit, and loss (Mises, 1920, 1949). A central algorithm can assign shadow prices, but shadow prices are not identical to prices generated by rivalrous bidding among owners who bear the consequences of their decisions. The Mises-Hayek argument is therefore twofold: Mises emphasizes the indispensability of real prices for capital calculation, while Hayek emphasizes the role of prices in communicating dispersed knowledge (Mises, 1920; Hayek, 1945). Contemporary interpretations sometimes separate these arguments too sharply, but they are best understood as complementary. Without market institutions, prices lose their grounding; without dispersed knowledge, prices lose their explanatory significance as social signals. This is why Lange's market-socialist response, however brilliant, did not fully answer the Austrian critique: it treated the price system too much as a computational device and too little as an institutional process of rivalry, error correction, and discovery (Lange, 1936, 1937; Lavoie, 1985a; Boettke, 2001).

The political dimension is equally important. A comprehensive AI-planning system would require immense data collection. To know what people want, what they can do, what they might produce, how they might consume, where they might move, and how they might respond to policy, the planner would be tempted to observe behavior at unprecedented scale. The dream of perfect planning therefore easily becomes the dream of perfect surveillance. This is not an accidental danger. It follows from the logic of the project. If the problem is that knowledge is dispersed in persons, then the central planner must either leave persons free to act on that knowledge or attempt to extract it from them. AI makes extraction more technically feasible, but not more morally legitimate. Hayek's political warning in *The Road to Serfdom* remains relevant here: comprehensive economic planning tends to require coercive political authority because individuals disagree about ends, and a unified economic plan must ultimately impose some ranking of those ends (Hayek, 1944). A society that tries to solve the knowledge problem through total data capture may destroy the liberal conditions under which genuine knowledge, creativity, dissent, and experimentation flourish. Recent concerns about surveillance capitalism and algorithmic power give this old Hayekian warning a new technological form (Zuboff, 2019).

This paper therefore advances a middle but firm thesis. Artificial intelligence makes Hayek more complicated, but not obsolete. AI expands the domain of feasible organizational planning; it does not abolish the knowledge problem. AI can centralize some decisions; it does not justify comprehensive economic centralization. AI can improve forecasts; it cannot replace entrepreneurial discovery. AI can process signals; it cannot create the institutional meaning of market prices. AI can infer preferences; it cannot fully represent human purposes. AI can empower governments and corporations; it can also threaten freedom if knowledge

becomes monopolized by centralized authorities. The lesson is not that Hayekians should reject AI. The lesson is that AI should be understood within a broader institutional theory of knowledge, prices, competition, and power.

The proper conclusion is therefore not anti-AI. On the contrary, a Hayekian society should use AI widely, but it should use AI polycentrically. AI should be distributed among individuals, firms, entrepreneurs, researchers, consumers, local governments, and civil society organizations. It should augment decentralized decision-making rather than absorb all decisions into a single administrative intelligence. The most productive future is not a return to old *laissez-faire* without technology, nor a new algorithmic socialism ruled by a central model. It is an institutional order in which AI strengthens the capacity of many actors to use their own knowledge, test their own plans, and coordinate through prices, competition, and voluntary exchange. Such a system would combine technological intelligence with institutional humility: it would recognize that intelligence is not only located in machines, experts, or states, but also in the dispersed practices, experiments, failures, and discoveries of free persons.

The rest of the paper develops this argument in detail. Section 2 reconstructs Hayek’s original theory of dispersed knowledge and the price system. Section 3 explains why AI creates a serious challenge to Hayekian economics by expanding the codification, aggregation, and prediction of knowledge. Section 4 distinguishes data from knowledge and argues that many economically decisive judgments remain tacit, contextual, and action-dependent. Section 5 explains why AI prediction cannot replace entrepreneurial discovery and novelty. Section 6 revisits Mises and the problem of economic calculation, emphasizing why simulated prices are not equivalent to market prices. Section 7 examines incentive problems in information revelation and the danger of algorithmic manipulation. Section 8 discusses surveillance, power, and the political economy of AI planning. Section 9 analyzes firms as islands of planning within markets and explains why AI may increase centralization inside capitalism without abolishing the need for market discipline. Section 10 concludes that Hayek’s central insight survives precisely because the age of artificial intelligence makes the organization of knowledge more—not less—important.

2 Hayek’s Original Argument: Dispersed Knowledge, Prices, and the Coordination of Plans

Hayek’s “The Use of Knowledge in Society” must be read not as a narrow article about information transmission, but as one of the deepest statements in modern social theory about the limits of central reason. Its central question is deceptively simple: what is the economic problem society must solve? The conventional textbook answer is that society must allocate scarce means among competing ends. Hayek did not deny this formulation, but he argued that it is incomplete and potentially misleading. If all ends, means, preferences, technologies, and resource constraints were already known to a single mind, then the allocation problem

would indeed be a problem of computation or engineering. One could imagine a planner solving a vast system of equations. But in actual society, the relevant facts are not given to a single mind. They are dispersed among individuals, firms, households, workers, consumers, entrepreneurs, scientists, artisans, farmers, financiers, and local officials. The first problem is therefore not how to compute an allocation from given data, but how to mobilize knowledge that exists only in fragmented, situated, and often inarticulate form (Hayek, 1937, 1945).

This was Hayek's great transformation of the socialist calculation debate. Mises had already argued that socialism could not perform rational economic calculation without private ownership of the means of production and market prices for capital goods (Mises, 1920, 1922). Hayek accepted this argument but pushed it further. The socialist problem was not merely that central planners lacked a correct set of prices. It was that no central planning authority could possess, in usable form, the constantly changing knowledge needed to set such prices and direct production. Hayek's critique therefore moved beyond the arithmetical question of whether a planning board could solve enough equations. The more basic question was where the relevant knowledge comes from, how it is discovered, how it is communicated, and how individuals are induced to act upon it (Hayek, 1935, 1940, 1945). In this sense, Hayek turned economics from a theory of allocation under given information into a theory of coordination under dispersed and incomplete knowledge.

The most famous phrase in Hayek's essay is his description of "the knowledge of the particular circumstances of time and place" (Hayek, 1945). This phrase is decisive. Hayek was not primarily concerned with scientific knowledge in the narrow sense. He did not deny the importance of technical expertise, engineering science, statistical analysis, or general economic principles. But he insisted that such knowledge is only one part of the total knowledge used in economic life. The other part is local, concrete, and contextual. A shopkeeper may know that a certain product sells only when placed near the entrance. A farmer may know that a particular field will not produce well this season despite official crop estimates. A worker may know that a machine is about to fail because of a subtle change in sound. A builder may know which subcontractor can be trusted. A merchant may know that customers in one neighborhood have changed their habits. A consumer may know that a product has become less desirable for reasons no survey yet captures. Such knowledge is not necessarily theoretical. It is not always statistical. It is frequently tacit, practical, and temporary. Yet it is economically decisive.

This local knowledge is also perishable. Its value often depends on being used quickly. A temporary shortage, an unexpected demand shift, a sudden technological opportunity, a change in weather, a new transport delay, a local labor problem, or a new consumer fashion may matter only for days, hours, or even minutes. The center may eventually learn about it, but by then the opportunity may have passed. This is why Hayek emphasizes not only the dispersion of knowledge but also its timeliness. The problem is not merely that knowledge is scattered; it is that it is scattered in forms that require immediate judgment and adaptation. A system that waits for all knowledge to be transmitted upward, processed centrally, and translated into commands will necessarily lose much of the value of that knowledge. The

“man on the spot” has an advantage not because he knows everything, but because he knows something relevant at the moment when it can still be used (Hayek, 1945).

Hayek’s analysis also implies that much economic knowledge cannot be separated from action. People often discover what they know only when they must decide. A buyer discovers the intensity of his preference when he must pay. An entrepreneur discovers whether an idea is viable when he risks capital. A worker discovers his comparative advantage through experience in different tasks. A firm discovers whether a production method works by using it under competitive pressure. Knowledge in this sense is not merely information stored in the mind. It is revealed, refined, and sometimes created through choice. This is why competition later becomes, in Hayek’s language, a “discovery procedure” (Hayek, 1948, 1968). Competition is not only a mechanism for minimizing costs once all facts are known. It is a procedure for discovering facts that would otherwise remain unknown.

The price system is Hayek’s answer to this problem. Prices allow dispersed individuals to coordinate their plans without requiring centralized possession of all relevant knowledge. The price of a good summarizes countless conditions: relative scarcity, consumer demand, alternative uses, production costs, expectations, risks, and competing opportunities. No single participant needs to understand the whole causal chain behind a price change. If the price of tin rises, users of tin need not know whether the cause is a mine collapse, a war, a new industrial use, transport disruption, or increased demand elsewhere. They only need to know that tin has become more valuable relative to other goods. The higher price induces them to economize, substitute, innovate, delay, recycle, or search for alternatives. Through such adjustments, the economy responds to facts that most participants do not know directly (Hayek, 1945).

This is why Hayek calls the price system a mechanism for communicating information. But “communication” here must be understood carefully. A price does not communicate a full report. It communicates just enough for decentralized actors to adjust their plans. It economizes on knowledge by reducing the amount each person must know. In a modern economy, no one needs to understand the entire structure of production. The consumer does not need to know the supply conditions of steel, copper, lithium, oil, microchips, shipping containers, and skilled labor. The producer does not need to know the full hierarchy of consumer preferences in every market. The investor does not need to know every technical detail of every enterprise. Each actor responds to a limited set of prices, profits, losses, wages, rents, and interest rates. The result is not perfect knowledge, but mutual adjustment under radical cognitive limits (Hayek, 1945).

This is also why Hayek’s theory is not merely an efficiency argument. It is an argument about the division of knowledge in society. Adam Smith had emphasized the division of labor; Hayek extended the logic to the division of knowledge. In a developed economy, knowledge is not concentrated but divided. The social problem is to coordinate this divided knowledge without destroying its local usefulness. Central planning attempts to unify decision-making by bringing knowledge to the center. The market order instead coordinates decentralized decision-making by allowing people to act on their own knowledge while using prices as

condensed signals from the wider economy. This is the institutional genius of the price system. It permits a social order more intelligent than any of its individual participants.

The point also explains why Hayek rejects the idea that the choice is between planning and no planning. There is always planning. Households plan their consumption, firms plan their production, workers plan their careers, investors plan portfolios, and entrepreneurs plan ventures. The real question is whether planning should be centralized in one comprehensive plan or distributed among many separate plans coordinated through prices and competition. Hayek explicitly states that the dispute is not whether planning should occur, but whether planning should be done centrally for the whole economic system or divided among many individuals (Hayek, 1945). This formulation remains essential in the AI age. Artificial intelligence does not eliminate planning; it changes who can plan, with what tools, at what scale, and under what institutional discipline.

Hayek's argument therefore differs sharply from the neoclassical image of perfect competition as a state of equilibrium in which all agents already know all relevant prices and opportunities. In that static model, competition appears almost redundant because the relevant information is already assumed. Hayek criticized this tendency in "Economics and Knowledge" and "The Meaning of Competition." Equilibrium theory, he argued, often assumes away the very process by which knowledge becomes coordinated (Hayek, 1937, 1948). Real competition is valuable precisely because people do not already know the best methods, lowest costs, true demands, or most valuable uses of resources. Competition is a process in which rival plans are tested against reality. It reveals who can produce more cheaply, which products consumers actually value, which technologies work, and which expectations are mistaken. To describe competition only as an equilibrium condition is therefore to miss its main epistemic function.

This insight connects Hayek to later Austrian theories of entrepreneurship. Kirzner's entrepreneur is not merely a manager who solves a known maximization problem. He is alert to unnoticed opportunities. Profit arises because something has been overlooked: a price difference, an unmet demand, an unused resource, a new combination, or a possibility others failed to see (Kirzner, 1973, 1985). Schumpeter's entrepreneur, though different in emphasis, also reveals the inadequacy of static allocation theory. Economic development proceeds through new combinations that disturb equilibrium and create new possibilities (Schumpeter, 1911, 1942). Hayek's knowledge argument provides the social-theoretic foundation for both views. A decentralized market order is superior not simply because it allocates known goods efficiently, but because it permits discovery under uncertainty.

This discovery function is precisely what comprehensive planning has difficulty reproducing. A central planner may collect statistics on existing production and consumption. But statistics mostly describe what has already happened, or what is already visible. They do not easily reveal what might happen if individuals were free to try new things. A ministry may know last year's output, last month's sales, or yesterday's inventories. It may not know the new product that a frustrated engineer might design, the new use that a consumer might discover, the new method that a small firm might test, or the new preference that

might emerge in response to cultural change. The market order elicits these possibilities through decentralized trial and error. Central planning tends to suppress them because it must decide in advance which experiments deserve resources.

The issue is not that central planners are foolish. Hayek's argument is not anti-intellectual in that sense. The problem is institutional. Even intelligent planners cannot know what can only be discovered through decentralized action. They must rely on reports, statistics, models, and administrative channels. These forms of knowledge are useful, but they are not identical to the living knowledge of economic actors. Moreover, the more centralized the system becomes, the more filtered the information reaching the center will be. Local actors may conceal bad news, exaggerate needs, understate capacity, or manipulate reports. Thus the knowledge problem is inseparable from the incentive problem. Information does not simply flow upward truthfully because the planner wants it. Institutions shape what people reveal, hide, distort, and create (Hayek, 1945; Lavoie, 1985a; Kornai, 1980, 1992).

The Misesian calculation argument reinforces this Hayekian point. Market prices for capital goods are not merely useful numbers; they are produced by a social process in which owners bid, exchange, risk, profit, and lose. Without private property in the means of production, there can be no genuine market for capital goods; without such markets, there can be no genuine prices reflecting opportunity costs; and without such prices, planners cannot rationally compare alternative uses of heterogeneous capital (Mises, 1920, 1949). Hayek adds that these prices are not important only because they permit arithmetic calculation. They are important because they condense dispersed knowledge into signals that guide action. Mises and Hayek therefore offer complementary arguments. Mises shows why calculation requires market prices; Hayek shows why market prices require decentralized knowledge and adjustment.

The market-socialist response associated with Oskar Lange attempted to answer Mises and Hayek by proposing that a central planning board could set prices, observe shortages and surpluses, and adjust prices by trial and error until supply equaled demand (Lange, 1936, 1937). This response is important because it anticipates many modern AI-planning proposals. If the planner has enough information and can adjust prices rapidly enough, perhaps planning can mimic the market. Hayek's reply, however, remains powerful. Trial-and-error pricing by a planning board treats the problem as if the relevant knowledge were contained in aggregate excess demand functions. But actual market competition does more than clear markets. It discovers costs, products, techniques, substitutes, qualities, and opportunities that no planning board can specify in advance. Lavoie's reinterpretation of the socialist calculation debate rightly emphasizes that the Austrian critique was not defeated by formal equilibrium models of market socialism because the real issue was rivalry, discovery, and institutional process (Lavoie, 1985a).

This also explains why shortages became a chronic feature of socialist economies. Kornai's analysis of the socialist system showed that shortage was not an accidental malfunction but a systemic consequence of soft budget constraints, administrative allocation, and distorted incentives (Kornai, 1980, 1992). Planners could collect enormous amounts of information, but

the information was filtered through bureaucratic incentives. Enterprises had reasons to hoard inputs, exaggerate needs, avoid risk, and fulfill plan indicators rather than satisfy consumers. The result was not merely computational inefficiency but institutionalized miscoordination. Hayek's knowledge problem therefore operated not as an abstract philosophical difficulty but as a practical economic pathology.

The relevance of Hayek's original argument to the AI age should now be clear. Hayek does not prove that all centralization is impossible or that all planning is irrational. Firms plan. Households plan. Governments plan. Modern economies require organizations, standards, infrastructure, and public institutions. Hayek's point is instead that the scope and success of planning depend on whether the planner possesses or can legitimately obtain the relevant knowledge, whether the knowledge can be articulated, whether incentives induce truthful revelation, whether errors are corrected, and whether decentralized actors remain free to discover better alternatives. The case for markets is therefore not a dogmatic rejection of intelligence or organization. It is a case for an institutional order that uses knowledge more effectively than any centralized mind can.

For this reason, Hayek's "The Use of Knowledge in Society" survives as the indispensable starting point for evaluating artificial intelligence and central planning. AI changes the technological capacity to collect, process, and predict information. It does not abolish the division of knowledge in society. The essential Hayekian question remains: does a given institutional arrangement enable people to use knowledge that is local, tacit, changing, and often not yet known until acted upon? If AI is used to strengthen the ability of many actors to perceive opportunities, compare alternatives, and coordinate through prices, it may deepen the Hayekian order. If AI is used to absorb decision-making into a centralized apparatus that suppresses prices, rivalry, and entrepreneurial discovery, it will repeat the old planning error in a technologically more advanced form. The machine may be new; the knowledge problem remains.

3 The Artificial Intelligence Challenge: Codification, Prediction, Centralization, and the New Case for Planning

The argument of the previous section reconstructed Hayek's original theory of dispersed knowledge, prices, and decentralized coordination. But if this paper is to be serious, it cannot simply restate Hayek as if the technological world of 1945 still existed. Artificial intelligence changes the empirical setting in which Hayek's argument must be evaluated. Hayek wrote before modern computers, before the internet, before cloud computing, before digital platforms, before machine learning, before real-time supply-chain software, before ubiquitous sensors, and before the massive digitization of economic life. The informational world in which his argument was formulated was one of price signals, paper records, fragmented local reports, slow statistical bureaus, and bureaucratic ministries whose capacity to collect, process, and interpret information was sharply limited. Today, by contrast, a vast and

growing portion of economic life is recorded in digital form. Purchases, searches, clicks, locations, messages, inventories, payments, shipments, machine operations, worker schedules, energy use, and consumer reviews can be collected and analyzed at enormous scale. It is therefore no longer sufficient to say that planners cannot process enough information. Artificial intelligence dramatically expands the domain of processable information. The serious question is whether this expansion merely improves particular forms of organization or whether it weakens the Hayekian case for decentralized markets at its root.

The strongest pro-AI argument begins from the observation that the boundary between tacit and explicit knowledge is not fixed. Hayek emphasized knowledge that is local, contextual, and often difficult to communicate to a central authority (Hayek, 1945). Michael Polanyi later gave philosophical depth to this problem by arguing that “we can know more than we can tell” (Polanyi, 1966). Much practical knowledge is embodied in skill, habit, perception, and judgment rather than stated in rules. For much of the twentieth century, this appeared to place a hard limit on central planning and automation alike. If people could not fully articulate how they recognized a face, repaired a machine, diagnosed a patient, drove a car, managed a store, or anticipated a customer’s need, then neither bureaucrats nor computers could easily reproduce such knowledge. David Autor called this “Polanyi’s paradox”: we know many things tacitly that we cannot explicitly explain, making them difficult to automate through traditional rule-based programming (Autor, 2014).

Modern AI weakens this version of the tacit-knowledge argument. Machine learning does not require all knowledge to be expressed in explicit rules. It can learn from examples, patterns, images, speech, text, sensor signals, and behavior. A person may not be able to state the rules by which she recognizes fraud, detects customer dissatisfaction, translates idioms, drives through traffic, or judges product quality, yet an AI system may learn to approximate these judgments from large datasets. In this sense, AI can codify some forms of knowledge without first requiring human beings to articulate them. It can convert repeated human behavior into models, routines, classifications, predictions, and recommendations. Knowledge that once seemed inseparable from the “man on the spot” may become increasingly transferable to software systems operating far from the local context.

This is the central insight of Brynjolfsson and Hitzig’s “AI’s Use of Knowledge in Society.” Their title deliberately reverses Hayek’s formulation and asks whether artificial intelligence changes the social use of knowledge. They argue that AI may codify local and tacit knowledge that previously favored decentralized decision-making. If a system can learn from worker behavior, customer interactions, managerial decisions, production histories, and digital traces, then some knowledge once possessed only by local actors can be extracted, centralized, and redeployed elsewhere. This may reduce the informational advantage of local managers and increase the efficient scale of hierarchical control. Brynjolfsson and Hitzig therefore predict that AI can shift the optimal locus of decision-making upward, increase firm size, raise industry concentration, and reduce the need for local managerial autonomy in some sectors (Brynjolfsson and Hitzig, 2025). Their argument does not prove that comprehensive socialist planning becomes viable, but it does create a powerful challenge

to a simple reading of Hayek. If AI can convert tacit knowledge into transferable knowledge, then one of Hayek's central premises becomes less absolute.

The second pro-AI argument concerns prediction. Hayek emphasized the inability of central planners to know rapidly changing local conditions (Hayek, 1945). But modern predictive systems can process real-time data on a scale and speed that no earlier planning apparatus could approach. Retailers can forecast demand using point-of-sale data, weather, online searches, seasonal patterns, local events, and competitor behavior. Logistics firms can reroute shipments by integrating traffic data, port congestion, weather conditions, fuel prices, and warehouse capacity. Manufacturers can forecast equipment failure from sensor data before human operators notice a problem. Hospitals can predict patient flows. Energy systems can forecast demand and manage grid loads. Financial institutions can detect suspicious transactions in milliseconds. These examples show that AI does not merely collect information; it generates forward-looking estimates that can guide action before shortages or failures become visible to human agents.

This predictive capacity matters because one of the traditional arguments for markets is that they allow fast adaptation to changing circumstances. If prices rise, decentralized agents adjust. But if AI systems can anticipate changes before prices fully move, then centralized or semi-centralized organizations may sometimes adjust faster than markets. A platform that sees search behavior, inventory levels, and supplier conditions may anticipate a demand surge before local retailers notice it. A central logistics system may detect a bottleneck before prices reveal it. A headquarters using AI may allocate inventory across regions more efficiently than store managers acting only on local knowledge. In such cases, prediction partly substitutes for ex post adjustment. The planner need not wait for price signals if it can infer future scarcity directly from data.

The third pro-AI argument concerns aggregation. Hayek's central planner failed because knowledge was dispersed and difficult to combine. But AI is precisely a technology of aggregation. It can combine structured and unstructured data from many sources: transactions, texts, images, voice recordings, contracts, satellite images, warehouse scanners, online behavior, industrial sensors, and social networks. Earlier bureaucracies faced severe limits in collecting and interpreting such data. Modern platforms already aggregate knowledge on a massive scale. Amazon, Walmart, Alibaba, Google, Microsoft, Apple, Meta, and other large digital firms coordinate economic activity using data infrastructures that would have been unimaginable to earlier planning theorists. These firms do not abolish markets, but they show that large-scale informational centralization is possible within a capitalist economy. The empirical existence of platform capitalism thus complicates any simple Hayekian claim that knowledge cannot be aggregated.

The fourth argument concerns internal planning within firms. As Coase argued, firms exist because some coordination can be performed more efficiently through hierarchy than through market exchange (Coase, 1937). Williamson later developed this into a theory of transaction-cost economizing, in which markets and hierarchies are alternative governance structures suited to different conditions (Williamson, 1975, 1985). AI changes these conditions. By

reducing the cost of monitoring, prediction, communication, and internal coordination, AI may enlarge the domain in which hierarchy outperforms market contracting. A firm that can observe operations in real time, standardize best practices, monitor performance, forecast demand, and allocate resources algorithmically can manage a larger and more complex organization than firms in earlier periods. AI therefore may push the Coasean boundary outward, increasing the efficient scope of the firm. In this sense, AI does not necessarily prove socialism, but it may prove more centralization inside capitalism.

The fifth argument concerns the relationship between AI and managerial discretion. In a pre-AI organization, local managers possessed information that headquarters could not easily see. They knew local workers, customers, suppliers, habits, and conditions. Headquarters had to delegate because it lacked detailed local knowledge. With AI, however, headquarters can increasingly observe local operations directly. Sales data, customer feedback, worker productivity, inventory flows, delivery times, machine performance, and local demand patterns can be monitored continuously. Algorithms can compare units, detect anomalies, recommend interventions, and impose standardized procedures. This reduces the need for local autonomy in some settings. Brynjolfsson and Hitzig’s prediction of reduced local managerial discretion follows from this logic (Brynjolfsson and Hitzig, 2025). If headquarters can know more and decide faster, then decentralization becomes less necessary.

The sixth argument concerns the possibility of algorithmic trial and error. Lange’s market-socialist model proposed that a central planning board could adjust prices in response to shortages and surpluses until markets cleared (Lange, 1936, 1937). Hayek and Mises rejected this as insufficient because it reduced markets to a computational equilibrium procedure and ignored discovery, incentives, and institutional conditions (Mises, 1920; Hayek, 1940, 1945; Lavoie, 1985a). But AI gives Lange’s argument new life. A modern planning board would not adjust prices slowly by bureaucratic decree. It could run simulations, use reinforcement learning, conduct experiments, infer demand elasticities, model substitution patterns, and update shadow prices continuously. It could operate more like a dynamic platform than a Soviet ministry. To the pro-AI planner, the old Langean trial-and-error model looks technologically premature rather than theoretically impossible. What Lange imagined abstractly, algorithmic systems might now implement more effectively.

The seventh argument concerns the increasing role of platforms as quasi-planning institutions. Digital platforms already coordinate vast numbers of decisions through algorithms rather than through simple market prices alone. Ride-sharing platforms match drivers and riders, set dynamic prices, allocate routes, and manage reputational systems. E-commerce platforms rank sellers, recommend products, manage logistics, enforce standards, and sometimes compete directly with third-party sellers. Food-delivery platforms coordinate restaurants, drivers, customers, and time-sensitive logistics. Cloud platforms allocate computing resources across users and regions. These systems are neither pure markets nor traditional firms. They are algorithmic governance structures. They use prices, but they also use rankings, recommendations, rules, search placement, access controls, and automated sanctions. The rise of platforms therefore suggests that modern capitalism is already moving

toward forms of AI-mediated coordination that blur the boundary between market and plan.

This platform evidence is important because it reveals a weakness in crude market-versus-planning dichotomies. In practice, modern economic coordination occurs through hybrids: firms inside markets, platforms inside capitalism, standards inside competition, state procurement inside private production, and algorithmic ranking inside decentralized exchange. AI strengthens these hybrid forms. A digital platform can be market-like because users choose, compete, enter, and exit; but it can also be planning-like because the platform designs the rules, controls visibility, sets fees, governs access, and shapes behavior. Such systems show that AI does not simply recreate Soviet planning. It creates new institutional forms that combine market incentives with centralized informational control. These forms require a more refined Hayekian analysis.

The eighth argument concerns public administration. Hayek's critique was directed primarily against comprehensive economic planning, not every form of state action. Still, AI may strengthen the case for more ambitious public coordination in domains where market failures are serious. Governments can use AI to improve infrastructure planning, traffic management, tax enforcement, fraud detection, disaster response, public health surveillance, energy-grid management, environmental monitoring, military logistics, and social-service delivery. In these domains, the relevant objective may be clearer than in comprehensive economic planning, and the data may be more measurable. For example, traffic systems can optimize flows; public-health agencies can forecast hospital demand; tax authorities can detect evasion; emergency services can allocate resources during disasters. These forms of AI-assisted administration may not refute Hayek, but they do show that greater planning capacity can produce real public benefits.

The ninth argument concerns industrial policy. If AI can improve forecasting, identify bottlenecks, map supply chains, and simulate technological trajectories, governments may become better at supporting strategic sectors. The older critique of industrial policy emphasized that governments lack knowledge of future winners and are vulnerable to lobbying and capture. AI cannot eliminate these problems, but it may reduce some informational limits. It can map dependencies in semiconductors, energy, pharmaceuticals, defense, critical minerals, and logistics. It can identify vulnerabilities, estimate spillovers, and evaluate alternative investments. In a world of geopolitical competition and supply-chain fragility, the pro-AI planner may argue that decentralized markets alone underprovide resilience, national security, and long-term technological capacity. AI therefore may strengthen not only firm-level planning but also state-guided development strategies.

The tenth argument concerns the scale of modern problems. Climate change, pandemics, financial contagion, cyber risk, demographic change, energy transition, and geopolitical supply-chain disruptions are systemic problems. They involve externalities, interdependence, and long time horizons. Market prices often fail to incorporate all relevant social costs and benefits. AI may help public authorities model such complex systems and coordinate responses. A climate transition, for instance, requires coordination across energy production, transport, construction, agriculture, finance, technology, and consumer behavior. A pandemic

requires coordination across hospitals, supply chains, testing, vaccines, communication, border controls, and behavioral responses. The pro-AI argument is that decentralized markets may be too fragmented to handle such systemic problems adequately, while AI-enhanced planning may allow more coherent responses.

The eleventh argument concerns welfare and personalization. A traditional objection to central planning is that planners cannot know individual preferences. But AI-based recommendation systems already personalize goods, services, information, entertainment, advertising, education, and even medical advice. They infer preferences from behavior and adapt in real time. A pro-AI planner may argue that the old image of uniform central planning is obsolete. Planning no longer needs to impose the same goods on everyone. Algorithmic systems can personalize allocation. They can recommend different diets, learning plans, medical treatments, job matches, transport routes, energy usage patterns, and consumption bundles. If AI can personalize at scale, then the contrast between centralized planning and individual preference becomes less sharp than in older socialist systems.

The twelfth argument concerns behavioral inefficiency. Markets assume that individuals reveal preferences through choices, but behavioral economics has shown that choices are often shaped by biases, framing effects, limited attention, addiction, misinformation, and present bias. AI systems might correct some of these failures by helping individuals choose better. A planner might claim that algorithmic guidance can reduce wasteful consumption, improve health, encourage saving, reduce energy use, and match workers to better opportunities. This argument does not belong to classical socialism, but it belongs to modern technocratic paternalism. If individuals do not always know or act upon their own best interests, then AI may appear to justify more active guidance by experts or public systems. This line of thought connects AI planning not only to socialist calculation but also to contemporary debates over behavioral public policy and paternalism (Rizzo and Whitman, 2019).

The thirteenth argument is that AI may reduce the epistemic advantage of competition itself. Competition discovers information because rival firms try different things. But if AI systems can simulate many possible strategies before actual experimentation, then some costly trial and error may be replaced by computational exploration. Drug discovery, materials science, product design, logistics optimization, and advertising already rely on simulation and machine learning. The pro-AI planner could argue that markets discover through wasteful rivalry, while AI discovers through faster, cheaper search. If a central or platform-based system can simulate alternatives, test them digitally, and deploy the best option, then competitive decentralization may become less necessary in some domains. This is a serious claim because it strikes at the Hayekian idea of competition as a discovery procedure.

Taken together, these arguments form a powerful case that artificial intelligence places real pressure on Hayekian political economy. AI can codify some tacit knowledge; predict demand and disruption; aggregate dispersed data; extend the feasible scale of hierarchical coordination; reduce the informational advantage of local managers; revive Langan trial-and-error planning; empower platforms as algorithmic governance structures; improve

public administration; strengthen industrial policy; address systemic risks; personalize allocation; correct behavioral failures; and replace some forms of market trial and error with computational simulation. No serious defense of Hayek can ignore these developments.

At the same time, it is crucial to state exactly what the pro-AI case proves and what it does not prove. It proves that the boundary between centralization and decentralization is technologically variable. It proves that some knowledge once considered tacit can be partially codified. It proves that firms and platforms may become larger and more centrally managed. It proves that AI can improve many forms of planning. It may even prove that some forms of state capacity and industrial strategy become more effective. But it does not yet prove that comprehensive central planning can outperform decentralized market coordination. It does not prove that data are equivalent to knowledge. It does not prove that prediction is equivalent to discovery. It does not prove that simulated prices are equivalent to market prices. It does not prove that individuals will truthfully reveal all relevant knowledge to a central system. It does not prove that innovation can be planned before it is discovered. And it does not prove that concentrated informational power is politically safe.

This distinction is the hinge of the entire paper. The AI challenge is real because it changes the technological conditions under which knowledge is collected and used. But Hayek's argument was never reducible to the claim that central planners lacked enough computing power. It was a broader institutional argument about dispersed knowledge, price signals, competition, incentives, discovery, and freedom. AI modifies several components of this argument, but it does not abolish the whole structure. The next sections therefore answer the pro-AI case not by denying the power of artificial intelligence, but by clarifying the limits of data, prediction, codification, simulation, and centralized control. The Hayekian rebuttal begins with the most basic distinction: data are not the same as knowledge.

4 Data Are Not Knowledge: Tacitness, Context, Judgment, and the Limits of Codification

The strongest pro-AI argument rests on a powerful but incomplete inference: because artificial intelligence can process vastly more data than earlier planners, it can therefore overcome the knowledge problem. The inference is tempting because modern economic life produces immense quantities of digital traces. Purchases, clicks, searches, locations, ratings, inventories, work schedules, production logs, machine readings, medical records, social-media behavior, and financial transactions all appear to make human action increasingly visible. If the problem facing earlier planners was informational scarcity, then a world of abundant data seems to alter the problem fundamentally. Yet this is precisely where a Hayekian analysis must insist on a conceptual distinction that is too often blurred: data are not the same as knowledge. Data are recorded signs, traces, measurements, categories, or observations. Knowledge involves interpretation, context, relevance, judgment, expectation, purpose, and action. The existence of more data does not automatically mean that the

economic knowledge needed for coordination has been captured.

Hayek’s original argument already contained this distinction, even though he did not formulate it in the language of contemporary data science. When he wrote that the “data” of the economic problem are not given to a single mind, he was not merely saying that the planner lacks a large enough database. He was objecting to the very assumption that the relevant data exist in a fully articulated, objective, transferable, and centrally usable form (Hayek, 1945). Much of what matters in economic life is known only by particular persons because they occupy particular positions. It is not simply a fact waiting to be uploaded; it is a situated understanding of a concrete circumstance. A seller knows whether demand is changing not only by observing sales numbers, but by interpreting customer questions, hesitation, mood, local gossip, seasonal routines, neighborhood shifts, and subtle differences in product reception. A worker knows that a machine is performing poorly not only from a sensor reading, but from long familiarity with its normal rhythm, sound, and behavior. An entrepreneur senses an opportunity not because a variable has crossed a threshold in a model, but because he interprets scattered signs as pointing toward a possible future not yet visible in aggregate statistics.

This point does not deny the value of data. On the contrary, data may be extremely valuable. Hayek himself did not reject statistics, expert knowledge, or general scientific knowledge. He rejected the belief that such knowledge exhausts what society needs to coordinate economic activity (Hayek, 1937, 1945). The question is not whether AI can improve the use of some data; clearly it can. The question is whether the most important economic knowledge can be reduced to data in a form that a central system can collect, interpret, and act upon better than decentralized agents. Here the answer is far less favorable to the planner. The danger of AI enthusiasm is that it mistakes legibility for knowledge. What is recorded becomes visible; what is visible becomes modeled; what is modeled becomes governable; and what is governable is then mistaken for the whole of social reality. But economic life includes many forms of knowledge that are not easily made legible.

Michael Polanyi’s theory of tacit knowledge is indispensable here. Polanyi argued that “we can know more than we can tell” (Polanyi, 1966). This does not mean merely that people have information they have not yet bothered to express. It means that many acts of knowing are embodied in skill, perception, attention, and practice. A pianist, surgeon, mechanic, trader, teacher, negotiator, or entrepreneur may act intelligently without being able to state all the rules governing that action. The knowledge is not fully separable from the performance. It is enacted. Traditional rule-based computing struggled with this because programmers had to specify explicit procedures. Modern machine learning changes the problem because it can learn from examples rather than rules. As Autor explains in his discussion of Polanyi’s paradox, tasks once difficult to automate because people could not state the rules may become more automatable when machines infer patterns from large datasets (Autor, 2014). This is a genuine breakthrough.

But it does not abolish tacit knowledge. It only shows that some tacit performance can be statistically approximated. The difference is crucial. An AI system may learn to classify

images, predict equipment failure, recommend products, or imitate expert judgments in a domain with abundant examples. Yet the tacit dimension of human action often involves not merely pattern recognition but responsibility, prudence, moral discrimination, and sensitivity to context. A physician does not merely classify symptoms; she decides how to speak to a frightened patient, whether to trust a self-report, when to deviate from a protocol, and how to weigh risks in a situation where statistical evidence is incomplete. A manager does not merely observe productivity metrics; he judges morale, loyalty, fatigue, resentment, creativity, and trust. An investor does not merely extrapolate past returns; she evaluates character, institutional change, political risk, and technological possibility. These judgments may use data, but they are not identical to data.

The problem becomes sharper when we recognize that data are produced through categories. A datum is never simply reality itself; it is reality selected, measured, classified, and recorded according to some prior scheme. Before AI can process the world, the world must be rendered into variables. But what counts as a variable? What is measured, and what is omitted? Which categories are used? Which time horizon matters? What is treated as noise? What is treated as success? These are not merely technical questions. They are judgments. A central planner using AI must decide what to count as relevant knowledge. But Hayek's argument is precisely that much relevance is discovered locally and unexpectedly. A fact that appears irrelevant to a central model may become decisive in a particular situation. A local custom, informal trust relation, seasonal habit, unrecorded skill, aesthetic preference, or entrepreneurial hunch may not enter the dataset. Yet these may determine whether a plan succeeds or fails.

This is one reason why centralized data systems often become overconfident. They see what they are built to see. The danger is not ignorance alone but false completeness. A planner with little information may know that he is ignorant. A planner with enormous data may believe that he sees the whole. This is especially dangerous because AI systems can produce outputs with an appearance of precision. Forecasts, rankings, risk scores, optimization recommendations, and simulations may look scientific even when their underlying categories omit crucial dimensions of reality. Hayek's critique of constructivist rationalism is relevant here: modern reason often overestimates its ability to design complex social orders because it confuses what can be articulated with what is known by society as a whole (Hayek, 1967, 1973). AI can intensify this rationalist temptation by making the social world appear more measurable than it actually is.

The issue of context is equally important. Information often changes meaning when removed from its setting. A decline in sales may indicate falling demand, but it may also indicate a temporary stockout, bad weather, a local festival, a competitor's short-term discount, a product-placement error, a change in staff behavior, or a shift in customer composition. A worker's lower output may indicate laziness, illness, burnout, defective tools, family crisis, poor management, or hidden experimentation. A consumer's failure to buy may indicate lack of interest, lack of income, waiting for payday, dislike of packaging, moral objection, or ignorance of the product's existence. AI can sometimes infer the cause by

combining many variables. But causal interpretation remains difficult because the same observable behavior may have different meanings in different contexts. Local actors may understand the meaning because they live within the situation. A central model may see only correlations.

Hayek's phrase "particular circumstances of time and place" should therefore be read not as a reference to small facts alone, but to context-dependent meaning (Hayek, 1945). The particularity matters because economic action is not performed in abstract space. It occurs in neighborhoods, firms, families, cultures, supply chains, legal systems, reputational networks, and historical moments. The same price, wage, product, or policy may produce different responses in different contexts. A central planner may know the average response but miss the local variation that matters for actual coordination. AI can improve the estimation of heterogeneity, but it cannot eliminate the problem of interpretation. It can cluster, classify, and predict, but the meaning of a situation may still be known best by those embedded in it.

There is also a temporal problem. Data are usually records of the past. They may be updated in real time, but even real-time data are traces of what has just happened. Knowledge, however, often concerns what might happen if one acts now. Entrepreneurial judgment is forward-looking in a stronger sense than statistical prediction. It involves imagining a future that does not yet have a data history. Kirzner describes entrepreneurship as alertness to opportunities that have not yet been noticed by others (Kirzner, 1973, 1985). Schumpeter emphasizes new combinations that disrupt existing routines (Schumpeter, 1911, 1942). In both cases, the decisive act is not merely prediction from existing data, but the creation or discovery of a new possibility. An AI system trained on past behavior may assist such discovery, but it cannot replace the human act of deciding that a new possibility is worth attempting despite the absence of confirmatory data.

This is why many important economic decisions are not reducible to optimization under known probabilities. They involve uncertainty in the Knightian sense: situations in which probabilities themselves are unknown or unstable (Knight, 1921). A central AI system can estimate risk when the relevant distribution is sufficiently stable. But entrepreneurs often operate where the distribution is not known. They must form judgments under ambiguity. Should a firm create a product no one has bought before? Should an investor back a technology whose market does not yet exist? Should a worker leave a stable job for a new occupation? Should a city invest in a neighborhood before others see its potential? These choices cannot be fully resolved by data because the relevant future depends partly on the choices themselves. Action creates evidence.

The distinction between data and knowledge is also visible in consumer preference. Modern platforms infer preferences from behavior, and in many cases they do this impressively. Recommendation systems can predict what books, films, songs, products, routes, or advertisements a person may prefer. But inferred preference is not the same as reflective preference. Human beings do not merely consume according to fixed utility functions. They learn, regret, aspire, imitate, mature, rebel, and transform themselves. A person may click

on what he later judges degrading. He may buy what he does not admire. He may fail to choose what would enrich him because he has not yet encountered it. Preferences are not merely revealed; they are cultivated and sometimes corrupted. AI can infer patterns of behavior, but it cannot settle the deeper question of what persons should value. This becomes dangerous when algorithmic systems treat predicted behavior as the definitive truth about human desire.

Hayek's argument does not require a romantic view of individuals as perfectly self-knowing. Individuals make mistakes. They are biased, impulsive, ignorant, and vulnerable to manipulation. But the fact that individuals are imperfect does not imply that a central system can know their good for them. The knowledge of ends is itself dispersed and plural. People disagree about what kind of life is worth living. Markets do not solve all moral disagreements, but they allow many plans to coexist without requiring a single hierarchy of ends. Central planning, especially AI-enhanced planning, risks replacing the plural search for ends with an optimized behavioral model. In that case, the planner no longer merely coordinates means; it begins to shape persons.

The problem of data quality further complicates the pro-AI case. AI systems depend on the data fed into them. But economic data are not neutral reflections of reality. They are affected by incentives. When people know that a system rewards certain indicators, they adapt their behavior to those indicators. Managers manipulate performance metrics. Firms optimize search rankings. Workers game productivity measures. Schools teach to tests. Hospitals code diagnoses strategically. Platforms are manipulated by bots, fake reviews, and coordinated campaigns. Goodhart's law states that when a measure becomes a target, it ceases to be a good measure. In an AI-planned economy, the central model's metrics would become targets on a massive scale. The result would not be perfect knowledge but strategic data production.

This is a direct extension of the socialist incentive problem analyzed by Kornai. In socialist economies, enterprises hoarded inputs, exaggerated needs, and fulfilled plan indicators rather than satisfying consumers (Kornai, 1980, 1992). The problem was not that planners had no data. They had enormous reporting systems. The problem was that the data were generated under distorted incentives. AI does not automatically fix this. It may detect some manipulation, but it also creates new incentives to manipulate the algorithm. If access to resources depends on what the model predicts or rewards, actors will learn to feed the model favorable signals. Centralized AI planning may therefore suffer from algorithmic versions of old socialist pathologies: metric gaming, hidden shortages, inflated reports, false compliance, and misrepresentation of local conditions.

Markets also contain manipulation, deception, advertising, herd behavior, and misinformation. Hayekians should not deny this. But competitive markets create multiple correction mechanisms. Consumers can exit. Rivals can expose defects. Prices can move. Losses can punish errors. New entrants can challenge incumbents. Investors can withdraw. Journalists, courts, auditors, and regulators may intervene. A centralized AI system, by contrast, may concentrate error. If the central model misclassifies a need, misprices an input, suppresses

an innovation, or rewards the wrong metric, the mistake can propagate through the whole system. The very scale that makes AI powerful also makes its errors systemic.

The question of responsibility is another reason data cannot be equated with knowledge. Knowledge in economic life is often tied to bearing consequences. A business owner who believes demand will rise must risk capital. A worker who believes he has a better opportunity must change jobs. An investor who believes a technology will succeed must commit funds. These actions reveal seriousness because they are costly. Market behavior therefore produces information not merely because people speak, but because they act under constraint. A survey response, preference profile, or digital trace may be cheap and shallow. A purchase, investment, resignation, invention, or entrepreneurial venture is more meaningful because it involves sacrifice and risk. Prices emerge from such costly commitments.

An AI planner may observe behavior, but if it removes the institutional context of property, competition, and profit-loss responsibility, it changes the meaning of the behavior it observes. Data generated in a market order cannot simply be transferred into a nonmarket planning order without loss. This is a crucial point. Much of the data that make AI powerful today are generated inside market or quasi-market environments: consumers choose, firms compete, sellers adjust, workers move, investors risk, and prices fluctuate. If comprehensive planning suppresses these processes, it may destroy the very data-generating mechanism on which its AI depends. The planner may inherit historical market data, but over time its own system would produce less reliable signals.

This is why shadow prices and simulated demand are not enough. A model can estimate what people might pay under market conditions, but willingness to pay is revealed most meaningfully when people actually face trade-offs. A consumer who says he values environmental quality may choose differently when prices rise. A firm that claims a project is efficient may act differently when its own capital is at risk. A worker who reports a preference for one job may change his mind after experiencing it. Knowledge emerges through real constraints. Simulation can assist, but it cannot fully replace the discipline of action.

The tacit dimension of knowledge also includes moral and relational knowledge. Trust, reputation, loyalty, resentment, credibility, and informal norms are economically significant. Contracts do not enforce themselves. Supply chains depend on trust. Teams depend on morale. Credit depends on reputation. Innovation depends on communities of practice. Much of this can be partially measured, but much remains qualitative and relational. A central system may know that a supplier has a high reliability score, but a local buyer may know that the supplier's owner is under stress, that a key employee has left, or that a promise made this time should not be trusted. A model may know that a worker's output is high, but a manager may know that the worker is poisoning team morale. These are not mystical facts. They are real economic knowledge whose meaning depends on human context.

The Hayekian point is not that AI has no access to such knowledge. Increasingly, AI may infer sentiment, reputation, trust networks, and organizational risk. But inference is

not equivalence. The more AI attempts to capture subtle human relations, the more it approaches surveillance and behavioral management. Thus the epistemic problem becomes a political problem. To collect more knowledge, the planner must observe more life. To observe more life, it must intrude deeper into privacy, association, speech, and conduct. At some point the attempt to solve the knowledge problem by data extraction begins to alter the society being known. People behave differently when they are constantly measured. Trust declines. Experimentation narrows. Dissent becomes risky. The result may be more data but less genuine knowledge.

This brings us to a central irony of AI planning. The more comprehensive the data system becomes, the more it may suppress the spontaneous sources of knowledge. People reveal knowledge not only because they are observed, but because they are free to try, fail, speak, buy, sell, move, imitate, and dissent. A liberal market order generates knowledge through freedom of action. Comprehensive AI planning attempts to replace freedom of action with observation and optimization. But observation is not a substitute for liberty. A society of monitored subjects may be more legible to the center while becoming less creative, less truthful, and less adaptive.

The conclusion of this section can be stated directly. AI greatly increases the ability to collect and process data. It can codify some tacit patterns, improve forecasts, detect correlations, and extend organizational control. But Hayek's knowledge problem concerns more than data scarcity. It concerns the social use of situated, tacit, contextual, forward-looking, incentive-laden, and action-dependent knowledge. Some of this knowledge can be digitized; much cannot. Some can be inferred; much must be discovered through decentralized experimentation. Some can be predicted; much is created by entrepreneurial action. Data are indispensable, but they are not the same as knowledge. The central planner's danger is to confuse the visible with the real, the measurable with the relevant, the predictable with the possible, and the modeled with the known.

For this reason, artificial intelligence does not abolish Hayek's argument. It clarifies it. Hayek's central insight was never that human beings lack enough calculating power. It was that social knowledge is divided among persons and becomes economically useful only through institutions that allow them to act on it. The price system, competition, property, profit and loss, and entrepreneurial freedom are not primitive substitutes for computation. They are social procedures for generating, testing, and communicating knowledge that cannot be fully centralized. AI can strengthen these procedures when it is placed in the hands of many actors. It can weaken them when it becomes an instrument of centralized control. The decisive issue is therefore not the intelligence of the machine, but the institutional order within which the machine is used.

5 Prediction Is Not Discovery: AI, Novelty, Entrepreneurship, and the Open Future

The second Hayekian rebuttal to AI planning follows naturally from the first. If data are not the same as knowledge, then prediction is not the same as discovery. Artificial intelligence is extraordinarily powerful when the future bears a sufficiently stable relationship to the past. Machine learning systems can infer patterns from historical and real-time data, estimate probabilities, identify correlations, classify cases, and recommend actions. In many domains, this is economically transformative. Forecasting demand, anticipating machine failure, detecting fraud, optimizing logistics, estimating credit risk, predicting hospital admissions, and recommending products all rely on the capacity to infer future outcomes from past and current patterns. But the market economy is not merely a prediction machine. It is also a discovery process. It does not only forecast future demand; it creates, reveals, and transforms demand. It does not only allocate known resources among known uses; it discovers new resources, new uses, new goods, new techniques, new forms of organization, and new human purposes. This is why AI prediction, however powerful, cannot replace the entrepreneurial and competitive discovery function emphasized by Hayek, Kirzner, Schumpeter, and Lavoie (Hayek, 1945, 1968; Kirzner, 1973, 1985; Schumpeter, 1911, 1942; Lavoie, 1985a).

Hayek's theory of competition is central here. In the textbook model of perfect competition, all relevant goods, prices, technologies, and preferences are already known. Competition then appears as an equilibrium condition rather than as an active process. Hayek objected that such a model assumes away the very reason competition matters. In "The Meaning of Competition," he argued that competition is valuable because it reveals facts that would otherwise remain unknown (Hayek, 1948). In his later essay "Competition as a Discovery Procedure," he made the point even more directly: competition is a method for discovering information that no one could possess before the competitive process unfolds (Hayek, 1968). This insight is indispensable in the age of AI. If the relevant knowledge already exists as data, then AI may help process it. But if the relevant knowledge does not yet exist because it will be generated by future experimentation, then no amount of prior data can fully substitute for the process that brings it into being.

The distinction can be put simply. Prediction asks: given what has happened, what is likely to happen next? Discovery asks: what possibilities exist that no one has yet recognized, tested, or valued? Prediction operates within a space of observed or inferable possibilities. Discovery enlarges the space itself. AI can search large possibility spaces, but the economic meaning of a possibility often emerges only through entrepreneurial action and social adoption. Before a new product exists, consumers may not know whether they want it. Before a new production method is tried, firms may not know whether it works. Before a new market is created, investors may not know whether it is profitable. Before a new lifestyle becomes available, individuals may not know whether they value it. The future is not merely hidden; it is partly made.

This point is especially important for understanding innovation. Schumpeter’s entrepreneur does not merely respond to existing demand. He introduces “new combinations”: new goods, new methods of production, new markets, new sources of supply, and new forms of organization (Schumpeter, 1911). Capitalism develops through creative destruction, the process by which new combinations disrupt old structures and transform economic life (Schumpeter, 1942). Such innovation cannot be fully planned from existing information because it changes the very information environment. A planner may know current demand for horses, carriages, candles, typewriters, telegraphs, landline phones, or film cameras. But that knowledge does not reveal the future demand for automobiles, electric lighting, computers, smartphones, internet platforms, or digital photography. The new good does not simply satisfy a pre-existing preference; it often educates, reshapes, or creates preference.

This is why Steve Jobs’s famous claim that consumers do not always know what they want until they are shown something new captures a deep economic truth. The statement is not a rejection of consumer sovereignty, but a recognition that preferences are often latent, undeveloped, and historically formed. People could not express demand for the iPhone in 1980 because the object, its ecosystem, and its associated practices did not exist. A central AI system trained on 1980 data might have predicted demand for better telephones, portable music players, cameras, maps, or computers. It could not have fully predicted the social form of smartphone life because that form emerged through design, experimentation, complementary innovation, platform building, and user adaptation. The same applies to search engines, social media, e-commerce, cloud computing, streaming, ride-sharing, mRNA vaccines, generative AI, and many other transformative technologies. Their value was not simply waiting in historical data. It emerged through the market process.

Kirzner’s theory of entrepreneurship provides a complementary account. For Kirzner, the entrepreneur is alert to unnoticed opportunities. Profit exists because markets are not in perfect equilibrium; something has been missed. There may be a price discrepancy, an unmet demand, an underused resource, an inefficient production method, or a possible exchange that others have failed to notice (Kirzner, 1973). Entrepreneurial discovery is therefore not identical to technological invention. It includes recognizing that a resource can be moved to a higher-valued use, that a consumer problem is unsolved, that a cost can be reduced, or that a product can be repositioned. This alertness cannot be fully centralized because it depends on the freedom of many individuals to notice different things from different positions. AI may assist alertness, but it cannot replace the decentralized multiplicity of perspectives through which opportunities are perceived.

The planner’s difficulty is not merely that there are too many possibilities to calculate. It is that no planner can know in advance which possibilities would be noticed by free agents under decentralized incentives. A central authority may employ brilliant scientists, engineers, economists, and AI systems. But even the most capable center must choose which experiments to fund, which ideas to approve, which risks to tolerate, and which failures to forgive. This requires judgment under uncertainty. In a market economy, many entrepreneurs make such judgments independently. Most fail. Some succeed. The successes

reveal information that could not have been known before the experiment. The failures also reveal information by closing unproductive paths. Decentralized trial and error is not a defect of capitalism; it is one of its epistemic virtues.

This argument is strengthened by Frank Knight's distinction between risk and uncertainty. Risk refers to situations in which probabilities are known or can be estimated from stable distributions. Uncertainty refers to situations in which the relevant probabilities are unknown or not meaningfully measurable (Knight, 1921). AI is extraordinarily powerful under risk when large datasets permit probabilistic estimation. But entrepreneurship often occurs under uncertainty. The entrepreneur does not merely calculate expected value from known probabilities. He forms a judgment about an uncertain future that may be transformed by his own action. Should one invest in a technology before demand exists? Should one create a new category of product? Should one enter a market others consider impossible? Should one abandon a dominant business model before the data clearly require it? These decisions are not reducible to prediction because the evidence becomes available only after action.

Lachmann's radical subjectivist economics is relevant here. Lachmann emphasized that capital is heterogeneous and that expectations differ among individuals (Lachmann, 1956, 1986). In a complex economy, different entrepreneurs hold different visions of the future. These visions cannot all be correct, but the rivalry among them is productive because it tests alternative interpretations of an uncertain world. Capital is not a homogeneous fund that can be allocated mechanically. It consists of complementary and specific goods whose value depends on plans. A machine, building, patent, software system, or worker skill may have very different value under different entrepreneurial visions. The market process tests these visions through profit and loss. A central AI planner may evaluate alternative plans, but it cannot reproduce the same pluralism of commitment if all plans must pass through one hierarchy of authorization.

The evolutionary economics of Nelson and Winter also supports this view. Economic change is not simply movement toward a known optimum. Firms possess routines, search locally, imitate, innovate, and adapt under selection pressures (Nelson and Winter, 1982). Innovation is uncertain, path-dependent, and cumulative. Many discoveries arise not from centralized optimization but from decentralized search. Different firms try different routines; some survive, others disappear. This evolutionary process is wasteful if judged against a fantasy of perfect foresight, but it is rational under ignorance. AI may improve search, but it cannot eliminate the need for variation and selection. Indeed, AI itself has advanced through decentralized experimentation among universities, firms, open-source communities, investors, users, and competing research paradigms. Its own history confirms the importance of dispersed discovery.

A pro-AI planner might respond that artificial intelligence can simulate experimentation. It can run millions of virtual tests, search chemical or design spaces, optimize logistics, generate product concepts, and compare strategies before resources are committed. This is true and important. AI can reduce the cost of experimentation. It can enlarge the range of possibilities considered. It can discover patterns that human beings miss. In scientific

and technological domains, AI-assisted search may become a major engine of innovation. But simulation is not the same as market discovery. A simulation requires a model of what matters. It requires objective functions, constraints, input variables, and assumptions about user behavior, technical feasibility, institutional conditions, and social adoption. The model may explore possibilities within its constructed space, but the market tests possibilities in the world. The world contains unknown unknowns: regulatory responses, cultural meanings, complementary innovations, consumer reinterpretations, competitor reactions, moral opposition, political shocks, and sheer accidents.

Moreover, the value of an innovation is often not purely technical. It depends on adoption. A technically superior product may fail because users dislike it, do not understand it, cannot integrate it into routines, distrust its producer, or prefer an inferior but familiar alternative. A technically modest product may succeed because it fits habits, identities, institutions, and complementary goods. AI simulation may help estimate these factors, but market adoption remains a discovery process. The planner cannot know the social meaning of an innovation simply by evaluating its technical properties. Many technologies become valuable only after users creatively appropriate them for purposes the inventor did not foresee. Markets allow such appropriation because they leave room for decentralized experimentation by users as well as producers.

The history of innovation is full of such surprises. Technologies invented for one purpose often become valuable for another. Firms discover unexpected customer segments. Users repurpose products. Complementary goods create new demand. Scientific tools become consumer goods; military technologies become civilian infrastructures; entertainment platforms become commercial marketplaces; communication tools become political instruments. A central planning system tends to classify technologies according to existing categories and official objectives. A decentralized market permits reinterpretation. This is one reason why the future cannot be fully predicted from the past. Novelty changes categories.

Hayek's price system is important in this discovery process because it provides feedback. When entrepreneurs try new combinations, prices, profits, losses, and market shares reveal whether resources are being moved toward higher-valued uses. These signals are imperfect and sometimes distorted, but they are continuous and decentralized. Profit is not merely a reward; it is evidence that an entrepreneur has discovered a way to create value under existing conditions. Loss is not merely punishment; it is evidence that a plan has failed to use resources in a sufficiently valuable way. Without profit and loss, experimentation loses discipline. A planner may fund many experiments, but unless failure has consequences and resources can be withdrawn, experimentation becomes bureaucratic patronage.

This is where Mises's calculation argument intersects with Schumpeterian and Hayekian discovery. Innovation requires not only imagination but also economic comparison. A new technology may be scientifically impressive but economically wasteful. A new product may be desirable but too costly relative to alternatives. A new infrastructure project may be politically attractive but less valuable than decentralized investments. Market prices allow innovators and investors to compare heterogeneous inputs and outputs. Without genuine

prices for capital goods, labor, land, intermediate inputs, and consumer goods, the planner cannot reliably determine whether an innovation economizes resources or merely absorbs them (Mises, 1920, 1949). AI can estimate costs and benefits, but those estimates depend on prices generated by real exchange. If the system suppresses markets, its innovation metrics become increasingly detached from opportunity cost.

The problem becomes more severe when innovation disrupts the plan itself. A central plan requires priorities. It allocates resources according to approved objectives. Radical innovation often threatens existing priorities, organizations, and political interests. In a market economy, an outsider can challenge incumbents by attracting consumers, workers, and investors. In a planned economy, the outsider must persuade the planning hierarchy. This creates a conservative bias. Even if the hierarchy uses AI to evaluate proposals, it may still reject innovations that disrupt official objectives, threaten politically favored sectors, or undermine existing bureaucratic authority. The problem is not intelligence but permission. Discovery requires permissionless entry to a degree that comprehensive planning cannot easily tolerate.

Lavoie's reinterpretation of the socialist calculation debate is important precisely because he emphasizes rivalry. The market is not merely a computational device for solving an allocation problem. It is a rivalrous process in which actors compete to interpret the future and discover better uses of resources (Lavoie, 1985a). This rivalry matters because no single interpretation should be allowed to monopolize the future. AI planning risks giving one model, or one hierarchy of models, excessive authority over which possibilities are worth pursuing. Even if the model is highly sophisticated, it remains a selection mechanism. Its outputs depend on training data, assumptions, objectives, and institutional incentives. If everyone must pass through it, the economy becomes epistemically narrower.

This narrowing danger is already visible in algorithmic platforms. Recommendation systems can help users discover goods, music, books, films, and ideas. But they can also channel attention toward what the model predicts will maximize engagement, sales, or retention. Such systems may reinforce existing preferences rather than broaden them. They may favor products similar to past successes. They may make it harder for radically new or difficult goods to be discovered. The same problem would be magnified in an AI-planned economy. If the central model predicts demand from past behavior, then goods that do not fit past behavior may be underproduced. Novelty requires some break with prediction. An economy ruled by prediction may become efficient at reproducing the known and weak at discovering the unknown.

This is one of the deepest dangers of confusing prediction with wisdom. Human development often depends on exposure to what we did not know we wanted or needed. Education, art, science, entrepreneurship, and moral life all involve surprise. A society organized too tightly around predicted preferences may become culturally and economically conservative in the name of efficiency. It may feed people what they already like, produce what models already validate, and discourage experiments whose value cannot yet be demonstrated. Hayek's defense of competition as discovery is therefore also a defense of

openness. The future should not be pre-authorized by the past.

The issue of preference formation is central. Standard welfare economics often treats preferences as given. But in historical reality, preferences are shaped by innovation, imitation, advertising, moral education, family life, religion, fashion, technology, and institutions. AI systems that infer preferences from behavior are therefore observing a moving and socially formed target. A person's "preference" for a product may reflect habit, addiction, lack of alternatives, manipulation, ignorance, or temporary mood. A society's demand pattern may reflect existing technological constraints rather than deep human ends. Entrepreneurs change preferences by introducing new possibilities. Artists change taste by creating new forms. Educators change aspiration by revealing higher goods. Political and moral movements change what people consider desirable. A central AI system that treats observed behavior as the truth of preference may freeze a historically contingent pattern into an administrative fact.

Hayek does not offer a full theory of preference formation, but his institutional argument is compatible with a plural and evolutionary view of human ends. Decentralized markets allow people to pursue different experiments in living, production, and consumption. They do not require one authority to decide which future preferences are legitimate. This is crucial because no AI system can determine, from data alone, whether a preference should be satisfied, discouraged, cultivated, or transformed. That is a moral and political judgment. The more comprehensive the AI plan, the more it must choose among competing visions of human flourishing. The market order, for all its imperfections, reduces the need for such centralized moral decision-making by allowing plural plans to coexist under general rules.

The problem of novelty also applies to labor and human capital. AI systems can match workers to jobs based on skills, experience, location, and predicted performance. This can improve labor-market efficiency. But human beings often discover their abilities through unexpected opportunities. A person may become a successful entrepreneur, scholar, artist, teacher, engineer, or manager only because a contingent path allowed talents to emerge. If an AI system assigns individuals according to predicted comparative advantage, it may improve short-run matching while limiting long-run self-discovery. Past performance is an imperfect guide to future vocation. A free labor market permits error, mobility, reinvention, and aspiration. A planned system may classify too early and too rigidly.

This is especially important in a dynamic economy. The most valuable skills of the future may not be visible today. Workers, firms, and educational institutions must experiment. Many forms of human capital are discovered through doing. AI can recommend training paths, but if such recommendations become authoritative, they may narrow human development. The Hayekian concern is not only that central planning misallocates goods, but that it misallocates persons by treating them as predictable units rather than as agents capable of growth and surprise.

AI also cannot eliminate the role of judgment in selecting among innovations. Suppose a central AI system produces thousands of promising technological possibilities. Which should receive resources? Which should be delayed? Which should be prohibited as

dangerous? Which should be allowed despite moral controversy? Which should be left to private experimentation? These decisions require institutional judgment. In a market order, such judgment is dispersed. Consumers, investors, entrepreneurs, regulators, courts, and communities all participate. In a central plan, judgment is concentrated. AI can inform the judgment, but it cannot make the legitimacy problem disappear.

Furthermore, innovation is deeply connected to dissent. New ideas often appear foolish, wasteful, immoral, or impossible to existing authorities. The history of science, industry, and culture is filled with outsiders who challenged accepted wisdom. A centralized AI system trained on accepted patterns may reproduce the epistemic conservatism of existing institutions. It may identify outliers as errors rather than pioneers. Decentralized markets and civil society provide spaces where dissenting projects can survive long enough to prove themselves. This is not guaranteed; markets also neglect many valuable ideas. But they provide multiple centers of judgment. That multiplicity is a defense against centralized error.

The AI planner might reply that a sufficiently advanced system could deliberately encourage diversity, experimentation, and novelty. It could allocate resources to high-risk projects, simulate alternative futures, and maintain portfolios of innovation. This is possible. Modern states and firms already do some of this. Research agencies, venture capital portfolios, corporate labs, and prize systems all allocate resources to uncertain projects. A Hayekian need not oppose such institutions categorically. The issue is whether they replace or supplement decentralized discovery. When AI-assisted planning funds experimentation within a larger market order, it may be beneficial. When it becomes the gatekeeper of all experimentation, it becomes dangerous. A portfolio selected by one authority cannot substitute for a society-wide plurality of experiments.

Indeed, the success of AI itself illustrates this argument. Artificial intelligence did not emerge from a single comprehensive plan. It developed through universities, government funding, private laboratories, open-source communities, venture capital, chip manufacturers, cloud providers, user communities, and global competition. Different paradigms rose and fell. Expert systems gave way to statistical learning; neural networks went through periods of enthusiasm and decline; transformer architectures emerged from specific research contexts and then were scaled through massive private and public investment. No central planner foresaw the full path. The development of AI was itself a Hayekian-Schumpeterian process of dispersed knowledge, rivalry, experimentation, and creative recombination.

This history should make AI triumphalism more modest. The machine that now seems to promise central intelligence was produced by decentralized intelligence. The data, chips, software, research, capital, and human talent behind AI were coordinated through markets, universities, open scientific exchange, state funding, entrepreneurial risk, and global competition. To use AI as an argument against decentralized discovery is therefore paradoxical. AI is not proof that planning can replace markets; it is evidence that markets and plural institutions can generate tools of extraordinary power.

The central conclusion of this section is that AI's predictive power, though profound,

does not overcome the problem of novelty. Prediction is strongest when the future resembles the past. Economic progress often occurs when the future breaks from the past. The price system, competition, entrepreneurship, and profit-loss discipline remain indispensable because they do not merely forecast within known possibilities; they generate and test new possibilities. AI can assist this process by widening search, improving information, lowering experimentation costs, and helping actors make better decisions. But AI becomes dangerous when it is imagined as a substitute for the decentralized discovery process itself.

Hayek's relevance therefore increases rather than decreases in a world of artificial intelligence. The more powerful prediction becomes, the more important it is to preserve institutions that allow the unpredicted to emerge. A society governed only by prediction risks becoming a society governed by yesterday's patterns. A market society, at its best, permits tomorrow's surprises. The task is not to reject AI, but to place AI within an open order of entrepreneurial discovery. AI should help many people experiment, not authorize one center to decide in advance which experiments matter. It should expand the range of possible futures, not narrow the future to what the model already expects.

6 Simulated Prices Are Not Market Prices: Mises, Hayek, Economic Calculation, and the Limits of Algorithmic Allocation

The third Hayekian rebuttal to AI planning must return to Mises. Hayek's knowledge problem and Mises's calculation problem are often treated as separate arguments, but for the purposes of evaluating artificial intelligence they must be understood together. Hayek explains why knowledge is dispersed and why the price system communicates fragments of that knowledge in usable form (Hayek, 1945). Mises explains why genuine economic calculation requires market prices for capital goods, and why such prices cannot exist without private ownership of the means of production (Mises, 1920, 1922, 1949). AI may process more data than any earlier planner; it may forecast shortages, simulate demand, and compute shadow prices. But if the institutional basis of market pricing is absent, the planner still faces a fundamental problem: it lacks genuine prices that reveal opportunity costs through exchange, rivalry, property, profit, and loss. Simulated prices may guide administrative decisions, but they are not the same as market prices. This distinction is central to the survival of the Mises-Hayek critique in the age of artificial intelligence.

Mises's original argument in "Economic Calculation in the Socialist Commonwealth" was not that socialist planners could not count physical quantities. A planning agency could count tons of steel, hours of labor, units of electricity, acres of land, kilometers of rail, hospital beds, tractors, machine tools, and houses. The problem was that economic production requires choosing among alternative uses of heterogeneous resources. Steel can be used for bridges, cars, railways, ships, tools, weapons, buildings, pipelines, and machines. Skilled labor can be allocated to medicine, engineering, teaching, programming,

construction, finance, research, art, or administration. Capital goods are not homogeneous lumps of “capital”; they are specific, complementary, durable, and embedded in production plans. Rational allocation requires comparing the value of alternative uses. For Mises, such comparison requires money prices generated in markets for capital goods (Mises, 1920, 1949).

This point is often misunderstood. Mises did not say that socialism is impossible because planners cannot know consumer needs at all. Nor did he say that physical planning cannot produce anything. A socialist economy can build factories, produce steel, run trains, and allocate labor by command. The calculation problem is more precise: without market prices for the means of production, planners cannot know whether one production plan uses scarce capital goods more economically than another. They may know that a bridge can be built using one combination of concrete, steel, labor, and machinery or another combination with more timber, more labor, and less steel. But they cannot rationally compare these heterogeneous input combinations without prices expressing relative opportunity costs. They may know technical feasibility, but not economic efficiency (Mises, 1920).

AI does not automatically solve this problem because the problem is not merely computational. A powerful algorithm can compare millions of technical input-output combinations. It can optimize under constraints. It can simulate resource flows. It can estimate shadow values. But the values used in the calculation must come from somewhere. If they are derived from actual market prices, then the AI system is relying on the market. If they are administratively assigned, then they require the planner to know relative opportunity costs without the institutional process that reveals them. If they are inferred from past market data, then the system depends on historical price information generated under conditions it may no longer preserve. If comprehensive planning abolishes or severely weakens private capital markets, then over time the informational foundation of those prices erodes.

This is why the claim that “AI can calculate prices” is ambiguous. AI can compute numbers that play the role of prices inside a model. It can generate shadow prices for constraints in an optimization problem. It can estimate willingness to pay from behavior. It can forecast demand curves and supply responses. It can adjust administrative prices in response to shortages and surpluses. But market prices are not simply numbers attached to goods. They are outcomes of exchange among agents who hold property rights, face budget constraints, possess different knowledge, and bear consequences. A price is a social fact produced by rival bids and offers. Its epistemic force comes not only from information aggregation but from costly commitment. People reveal valuations by giving up something else.

The opportunity cost of a resource is not an engineering property. It depends on all alternative uses and all competing valuations of that resource across society. A ton of copper has no single objective economic value independent of what people might do with it. Its value depends on uses in wiring, electronics, construction, vehicles, power grids, military systems, art, and countless other applications, as well as expectations about future scarcity and innovation. Market prices emerge because different users bid against one another under

scarcity. The resulting price is imperfect, but it reflects a real contest among alternative plans. A simulated price may estimate this contest, but without the contest itself, it lacks the same epistemic grounding.

Hayek's contribution deepens this point. In "The Use of Knowledge in Society," Hayek emphasizes that prices allow individuals to respond to changes whose causes they do not know (Hayek, 1945). If tin becomes scarce, users of tin need not know why. A higher price tells them to economize. But this price signal exists because users, suppliers, speculators, producers, and consumers are all acting on their own knowledge. The price is not imposed from above after the planner has understood all circumstances. It arises from decentralized adjustments. Its advantage is precisely that it economizes on the need for comprehensive understanding. An AI planner that tries to compute the "correct" tin price must collect and interpret the very knowledge that the market price system uses without centralizing.

Lange's market-socialist response remains the classic attempt to answer this problem. Lange argued that a central planning board could set prices for producer goods, instruct managers to minimize costs and produce where price equals marginal cost, observe shortages and surpluses, and then adjust prices by trial and error until equilibrium is reached (Lange, 1936, 1937). In this model, the planning board uses prices without private ownership of capital. The idea is elegant, and in the AI age it appears newly plausible. A modern algorithm could perform Langean trial and error much faster than any human planning board. It could update prices continuously, use reinforcement learning, simulate production networks, estimate elasticities, and respond to inventories in real time. Thus AI seems to revive Lange against Mises and Hayek.

But the apparent revival rests on a narrowing of the problem. Lange treated the price system mainly as a computational device for solving simultaneous equations or clearing markets. Mises and Hayek treated it as an institutional process. Lavoie's major contribution was to show that the socialist calculation debate was misunderstood when reduced to static equilibrium computation. The Austrian critique concerned rivalry, discovery, entrepreneurship, and the institutional generation of knowledge (Lavoie, 1985a). Trial-and-error pricing can adjust to observed shortages and surpluses among known goods, but it cannot fully reproduce the competitive process that discovers unknown costs, unknown demands, unknown substitutes, and unknown production methods. It can imitate one surface feature of markets while missing the deeper mechanism.

This distinction is essential for AI planning. Suppose an AI planning system observes a shortage of a certain input and raises its administrative price. Users economize, substitute, or request exemptions. The algorithm updates again. This resembles market adjustment. But in a genuine market, price changes also stimulate entrepreneurial search. Firms seek new suppliers, invent substitutes, redesign products, speculate on future scarcity, enter production, exit unprofitable lines, and bear gains or losses. The adjustment is not merely parametric. It is creative and institutional. A market price is embedded in rights to buy, sell, own, transform, and profit from resources. An administrative shadow price may instruct managers how to behave, but if managers do not own assets, cannot freely enter markets,

cannot capture profits, and do not bear losses, their response will differ.

The problem of incentives therefore cannot be separated from calculation. Mises emphasized ownership because owners have reasons to economize. They gain from moving resources to higher-valued uses and lose from wasting them (Mises, 1949). In a planned system, managers may be told to minimize costs, but their rewards often depend on plan fulfillment, political favor, risk avoidance, or bureaucratic compliance. Kornai showed that socialist enterprises operated under soft budget constraints, chronic shortage, and weak incentives to economize in the market sense (Kornai, 1980, 1992). Inputs were hoarded because future availability was uncertain. Targets were negotiated strategically. Quality was often sacrificed to quantity. Enterprises sought easier plans rather than consumer satisfaction. These were not merely errors in computation; they were systemic consequences of the institutional environment.

AI may detect some of these distortions, but it cannot abolish them by calculation alone. If managers know that the algorithm allocates resources based on reported need, they have incentives to exaggerate need. If the algorithm rewards output metrics, they optimize metrics. If the system punishes visible failure, they avoid risky innovation. If the model predicts future budgets from past utilization, departments spend resources to protect future allocations. These are algorithmic versions of old planning pathologies. A more intelligent center may monitor more closely, but monitoring itself creates further strategic behavior. The result can be a high-tech version of soft-budget politics rather than true economic calculation.

Market prices also incorporate time and expectations. An asset price reflects not only current use but expected future uses, risks, and opportunities. Interest rates coordinate intertemporal allocation by reflecting time preference, expected productivity, inflation expectations, risk, and monetary conditions. Commodity futures prices incorporate expectations about future scarcity. Equity prices reflect anticipated profits. Land prices reflect expected future development. These prices can be wrong, sometimes disastrously wrong, but their errors emerge from rival expectations backed by capital. A central AI system can forecast the future, but if investment decisions are not tested by decentralized ownership and profit-loss responsibility, the system may lack a comparable mechanism for disciplining expectations.

The issue of capital heterogeneity makes this problem even deeper. Austrian capital theory, especially in Mises and Lachmann, emphasizes that capital goods are not interchangeable units but heterogeneous components of plans (Mises, 1949; Lachmann, 1956). Their value depends on complementarity. A machine is valuable if combined with certain workers, materials, software, energy, logistics, and final demand. If the plan changes, the value of the machine changes. Central planning must therefore evaluate not only isolated goods but complex capital combinations over time. AI can model such networks better than previous planners, but the underlying valuation still depends on competing entrepreneurial plans. In a market economy, different entrepreneurs form different capital combinations and test them. In a planned economy, the center must choose among them administratively.

This is where Schumpeterian innovation reinforces Misesian calculation. New combi-

nations require capital reallocation (Schumpeter, 1911). But how does a planner know whether to move resources from an old industry to a new one? A new technology may appear promising but consume scarce engineers, chips, energy, and research funds that might have produced greater value elsewhere. Without market prices and capital markets, the planner lacks reliable measures of opportunity cost. AI may estimate social returns, but such estimates are necessarily model-dependent. In markets, venture capital, public equity, debt, retained earnings, and entrepreneurial ownership create many channels through which rival judgments about the future can be funded and tested. The capital market is not merely a financing mechanism; it is a decentralized selection system for plans.

A central AI planner could attempt to recreate this by running internal competitions, assigning experimental budgets, and evaluating outcomes. Some public agencies and large firms do this successfully in limited domains. But such internal quasi-markets remain different from open markets. They operate under administrative criteria. Entry is limited. Failure may be politicized. Projects depend on approval by committees or algorithms. The range of experiments is shaped by official objectives. Open capital markets, by contrast, allow outsiders to challenge dominant views. They permit eccentric, unfashionable, or disruptive visions to attract resources if someone is willing to bear the risk. This permissionless quality is central to discovery and calculation alike.

The problem of quality further illustrates why simulated prices are insufficient. Socialist systems often fulfilled quantitative targets while neglecting quality because plans specified measurable outputs (Kornai, 1980, 1992). AI can measure quality better than old planning systems through reviews, defect rates, sensor data, return rates, and user feedback. This is a real improvement. Yet quality is multidimensional and often subjective. Consumers care about durability, design, status, symbolism, ethical production, repairability, convenience, beauty, trust, and identity. Some of these can be measured; others emerge in use. Market prices for differentiated goods reflect consumers' willingness to pay for bundles of qualities. An administrative system may collect ratings, but ratings are not the same as purchase decisions under budget constraints. People may praise a product they would not buy at cost, or buy a product they would not praise publicly. Price-mediated choice remains a distinctive source of knowledge.

The same is true of labor. A planner may assign wages or labor credits based on skill, effort, need, or social priority. AI may estimate productivity more accurately than older systems. But wages in a market economy also coordinate mobility, training, occupational choice, regional movement, and entrepreneurial hiring. They reveal where labor is scarce relative to demand. They induce people to acquire skills and move to higher-valued uses. If wages are administratively compressed or politically determined, labor allocation becomes distorted. AI may recommend assignments, but individuals still possess private knowledge about their abilities, aspirations, family constraints, tolerance for risk, and willingness to sacrifice. A labor market allows such knowledge to be acted upon. A planning system must extract or override it.

The pro-AI planner might reply that modern systems need not abolish all markets. They

could preserve consumer choice and some labor mobility while centrally planning investment and major production. This is closer to market socialism or indicative planning than to full socialism. It is also more plausible. But to the extent such a system relies on genuine markets, it concedes the Mises-Hayek point. Its success depends on market prices generated outside the plan. The more it suppresses markets in capital goods, entrepreneurship, and investment, the more severe the calculation problem becomes. The more it preserves them, the less comprehensive its central planning is. AI may improve hybrid systems, but hybridity is not a refutation of the calculation problem; it is a response to it.

This is also the lesson of large capitalist firms. Firms internally plan. They use transfer prices, budgets, managerial hierarchies, performance indicators, and now AI systems. But they remain embedded in external markets. A firm can centrally plan production because it buys inputs at market prices, sells outputs to consumers, borrows or raises capital in financial markets, hires labor in labor markets, and faces competition. Internal planning is disciplined by external price signals. If a firm's internal AI system misallocates resources, losses eventually appear. If consumers dislike its products, revenues fall. If rivals do better, market share declines. If capital is wasted, investors react. Thus the success of planning inside firms does not show that comprehensive planning can dispense with markets. It shows that planning works best when surrounded by markets.

This Coasean point is central to understanding AI. AI may push the boundary of the firm outward by lowering the cost of internal coordination (Coase, 1937; Williamson, 1975, 1985; Brynjolfsson and Hitzig, 2025). But even very large firms depend on market prices. Amazon's logistics algorithms, Walmart's inventory systems, Toyota's production methods, Apple's supply-chain coordination, and cloud-computing platforms all rely on external markets for labor, components, capital, land, energy, and consumer demand. Their internal algorithms are powerful because they operate within a price environment. They do not prove that a whole society can be run as one firm. As Hayek warned, extending the logic of organization to the whole economy destroys the external market signals that organizations themselves need (Hayek, 1945).

One might object that a sufficiently advanced AI could estimate opportunity costs directly from physical constraints and observed preferences. But this returns us to the original problem. Preferences are not fixed objects waiting to be measured; they are revealed through choice and transformed through experience. Production possibilities are not fully known; they are discovered through entrepreneurial experimentation. Capital values depend on future plans; future plans are plural and uncertain. Opportunity cost is therefore not simply a physical shadow of scarcity. It is a social valuation produced by rival uses under conditions of freedom and constraint. AI can estimate it, but the estimation must be checked against real exchanges.

This does not mean that shadow prices are useless. They are valuable in firms, public agencies, cost-benefit analysis, environmental policy, infrastructure planning, and constrained optimization problems. A government may use shadow prices to evaluate carbon emissions, traffic congestion, public health interventions, or defense logistics. A firm may use internal

transfer prices to allocate resources among divisions. But shadow prices are most reliable when anchored in or disciplined by market prices. The danger arises when shadow prices are imagined to replace the entire market process. At that point they become administrative artifacts, vulnerable to model error, political pressure, and strategic manipulation.

The political problem is unavoidable. Market prices emerge from decentralized exchange. Administrative prices are set by authority. Even when an algorithm recommends them, some institution chooses the objective function, the constraints, the weights, and the permissible trade-offs. Should housing be cheap or land conserved? Should energy prices reflect climate costs, geopolitical security, consumer affordability, or industrial competitiveness? Should scarce medical resources be allocated by willingness to pay, need, expected benefit, age, productivity, or equality? These questions cannot be answered by calculation alone. They involve moral and political judgment. A central AI pricing system may hide such judgments behind technical language, but it cannot eliminate them.

In markets, moral and political questions do not disappear. Societies still regulate externalities, redistribute income, provide public goods, and restrict harmful exchanges. But the market order limits the range of decisions that must be made collectively. Many trade-offs are handled through decentralized choice. A comprehensive AI planning system expands the domain of collective decision and therefore intensifies political conflict. The planner must decide the relative importance of countless ends. Prices become instruments of policy rather than results of exchange. Calculation becomes politicized.

Hayek's broader warning against constructivist rationalism applies precisely here. Modern planners often believe that because they can describe an order in abstract terms, they can design and control it. But many social orders are the result of evolved rules and decentralized adaptation rather than deliberate construction (Hayek, 1967, 1973). The price system is such an order. It is not perfect, but its complexity exceeds deliberate design. AI increases the temptation to believe that what can be modeled can be governed. Yet modeling an order is not the same as replacing the order. The map may become richer, but it is still not the territory.

The central conclusion of this section is therefore clear. Artificial intelligence can calculate simulated prices, shadow prices, forecasts, and optimal allocations inside specified models. It can greatly improve administrative planning. It can make Langean trial-and-error pricing faster and more sophisticated. It can assist firms and governments in allocating resources. But it cannot, by computation alone, replace the institutional process through which market prices emerge. Genuine prices are produced by private property, exchange, competition, entrepreneurial judgment, profit, loss, and rivalrous bidding for scarce resources. They are not merely numerical summaries of data; they are social signals generated by action under constraint.

Mises and Hayek remain relevant because they understood that economic calculation is institutional before it is computational. Mises showed that without market prices for capital goods, rational comparison of heterogeneous production plans becomes impossible (Mises, 1920, 1949). Hayek showed that prices communicate dispersed knowledge that no central

authority can possess in full (Hayek, 1945). AI modifies the technical capacity of planners, but it does not abolish the need for market-generated opportunity costs. The planner may possess better machines than ever before; it still needs prices that no machine can generate in isolation from markets. Simulated prices may be useful servants. They are dangerous masters.

7 Incentives, Truthful Revelation, and Algorithmic Manipulation: Why Better Monitoring Does Not Solve the Knowledge Problem

The fourth Hayekian rebuttal to AI planning concerns incentives. Even if data could be distinguished from knowledge, even if prediction could be distinguished from discovery, and even if simulated prices could be distinguished from market prices, a central AI system would still face a further difficulty: knowledge does not reveal itself automatically. People do not simply transmit everything they know to the center in truthful, timely, and usable form. They reveal, conceal, distort, exaggerate, suppress, or strategically produce information depending on the incentives they face. The knowledge problem is therefore inseparable from the incentive problem. Hayek emphasized the dispersion of knowledge, but the institutional use of knowledge also depends on whether people have reason to act on what they know, bear the consequences of their actions, and reveal information in ways that others can rely upon (Hayek, 1945). Mises's calculation argument similarly presupposes that prices are meaningful because they arise from exchange under property, profit, and loss rather than from costless declarations of preference or administrative reports (Mises, 1920, 1949). Artificial intelligence can increase monitoring capacity, but it does not eliminate strategic behavior. Indeed, by making allocation depend more heavily on measured indicators, AI may create new and more sophisticated forms of manipulation.

This is one of the central lessons of the socialist experience. Planned economies did not fail merely because planners lacked enough statistics. They possessed enormous reporting systems. Ministries collected output data, labor data, input requirements, production targets, inventories, and fulfillment reports. The difficulty was that these data were generated within a system of distorted incentives. Enterprises had reasons to bargain for easier targets, exaggerate input needs, hide capacity, hoard materials, conceal bad news, and satisfy plan indicators rather than consumers. Kornai's analysis of shortage economies showed that chronic shortage was not an accidental technical malfunction but a systemic property of socialist coordination under soft budget constraints and administrative allocation (Kornai, 1980, 1992). The planner's knowledge was not simply incomplete; it was politically and strategically filtered.

AI does not remove this problem. A more advanced planning apparatus may monitor enterprises, workers, consumers, and local officials more closely. It may use anomaly detection, satellite images, sensors, financial traces, production logs, cameras, digital payments, and

predictive models to identify misreporting. This would make some forms of deception harder. But the monitored actors would adapt. If resources, promotions, punishments, permissions, or reputational scores depend on algorithmic measures, then individuals and organizations will learn to optimize those measures. A factory evaluated by output volume may sacrifice quality. A hospital evaluated by readmission rates may avoid high-risk patients. A school evaluated by test scores may teach to the test. A police department evaluated by crime statistics may manipulate reporting. A platform seller evaluated by ratings may purchase fake reviews. A worker evaluated by productivity software may maximize visible activity rather than useful work. The more consequential the metric, the stronger the incentive to game it.

This is the logic commonly associated with Goodhart's law: when a measure becomes a target, it ceases to be a good measure. Although the phrase is modern, the principle is deeply compatible with Hayek's and Kornai's critiques of planning. A central authority must simplify reality into indicators in order to govern it. But once those indicators determine rewards and punishments, they no longer merely measure behavior; they shape behavior. The data stream becomes endogenous to the control system. AI does not escape this problem because AI itself depends on measurable inputs. If the model rewards certain patterns, actors will produce those patterns. If the model penalizes certain anomalies, actors will hide them. If the model allocates resources based on projected need, organizations will manufacture evidence of need. The result is not transparent knowledge but strategic data production.

The old Soviet problem of plan fulfillment illustrates the general mechanism. If enterprises were required to produce a certain number of units, they might produce units of poor quality. If targets were set by weight, they might produce excessively heavy goods. If targets were set by value, they might produce high-priced goods regardless of consumer usefulness. These examples are often treated as anecdotes, but they express a deep institutional principle. Administrative indicators are substitutes for market feedback, and substitutes can be manipulated. Kornai's work showed that socialist firms responded rationally to the incentives embedded in the plan, even when their responses produced irrational social outcomes (Kornai, 1980, 1992). An AI-planned economy would face analogous problems. The indicators would be more numerous, dynamic, and sophisticated, but they would still become targets.

Nor is this problem limited to socialist economies. Modern capitalist organizations already display algorithmic gaming. Search engine optimization attempts to influence ranking algorithms. Social-media producers tailor content to recommendation systems. Online sellers manipulate reviews and ratings. Ride-share drivers and delivery workers learn how to respond strategically to platform incentives. Firms manage earnings reports. Universities optimize rankings. Hospitals adjust coding practices. Financial traders exploit predictable features of algorithms. These examples show that algorithmic governance does not eliminate strategic behavior; it creates new arenas for it. If such gaming occurs in competitive markets where exit and rivalry provide partial correction, it would be even

more consequential in a comprehensive planning system where the central algorithm governs access to resources.

This is why the incentive problem must be treated as epistemic, not merely moral. Dishonesty, gaming, and manipulation are not only ethical failures; they degrade knowledge. The planner's model learns from data produced by actors responding to the model. This creates feedback loops. If the model predicts demand based on subsidized consumption, it may mistake politically induced behavior for genuine preference. If it allocates investment based on reported productivity, it may reward manipulated output. If it assigns labor based on recorded performance, it may discourage unmeasured creativity. If it evaluates local governments by growth statistics, it may induce inflated reporting or wasteful construction. The AI system may become increasingly precise in analyzing increasingly corrupted data.

Hayek's defense of market prices is partly a response to this problem. Markets do not require individuals to reveal all their knowledge verbally or bureaucratically. They allow people to reveal information through costly action. A consumer reveals willingness to pay by actually paying. A producer reveals confidence by investing resources. An entrepreneur reveals belief in an opportunity by risking capital. A worker reveals preference and expectation by accepting, refusing, or leaving a job. An investor reveals judgment by buying or selling claims. These signals are not perfectly truthful, but they are disciplined by trade-offs. Cheap talk is replaced by costly choice. This is why prices convey more reliable information than surveys or administrative reports in many contexts (Hayek, 1945; Mises, 1949).

The key point is that market signals are incentive-compatible to a degree that planning reports are not. A person who claims to need more of a good may exaggerate if it is free or administratively allocated. But if he must pay a market price, his claim becomes more credible. A firm that claims its project is valuable may exaggerate if losses are socialized. But if investors must risk capital and bear losses, the claim is tested. A manager who claims to be efficient may manipulate metrics. But if consumers leave and competitors gain market share, the deception is harder to sustain. Profit and loss do not create perfect truth, but they impose consequences that make certain signals more meaningful.

AI planning may attempt to reproduce this through scoring systems, performance contracts, internal markets, and automated penalties. Such mechanisms can help, especially within firms and agencies. But they still depend on the authority defining the relevant metrics and enforcing the rules. A true market allows multiple evaluators: consumers, investors, suppliers, workers, competitors, lenders, insurers, and entrepreneurs. No single metric determines success. A firm can survive bad publicity if customers still value it; it can overcome short-term losses if investors believe in future profits; it can challenge incumbents through new products; it can pivot when one plan fails. A central algorithm, by contrast, tends to create a dominant evaluative structure. Even if the structure includes many variables, it remains a centralized criterion of allocation.

This distinction matters for innovation. If an AI-planned system rewards predictable performance, it may discourage experimentation. Novel projects often look inefficient at

first. They may generate losses, anomalies, delays, and uncertain outcomes. In a market, entrepreneurs and investors can choose to tolerate such uncertainty because they expect future gain. In a planned system, local actors may avoid risky projects if failure damages their algorithmic score. Officials may prefer safe compliance to uncertain innovation. Enterprises may produce what the model understands rather than what the future might value. This repeats Hayek's and Kirzner's point that discovery requires freedom to act on knowledge that is not yet socially validated (Hayek, 1968; Kirzner, 1973, 1985).

The problem is intensified by asymmetric information. Principal-agent theory shows that agents often possess information that principals cannot fully observe, and they may take actions that are difficult to monitor (Arrow, 1985). AI can reduce some information asymmetries by improving monitoring, but it cannot eliminate hidden action and hidden knowledge. A worker's effort, creativity, cooperation, and honesty cannot be fully captured by keystrokes, output counts, or surveillance. A manager's judgment cannot be fully captured by short-run metrics. A firm's true capability cannot be fully inferred from reports. A citizen's real preferences cannot be fully inferred from behavior under policy constraints. The agent always retains private knowledge, and the attempt to extract it may alter behavior.

Stiglitz's information economics also reinforces this point. Markets themselves suffer from asymmetric information, adverse selection, moral hazard, and incomplete contracts (Stiglitz, 2000). This means that Hayekian markets are not perfect. But the existence of information imperfections in markets does not imply that central planning solves them. A central authority faces even larger aggregation and incentive problems. It must gather information from many agents who understand the rules and have reasons to manipulate them. The relevant comparison is institutional: which system better elicits, tests, and corrects information under realistic incentives? Hayek's answer is that decentralized markets, though imperfect, possess superior mechanisms of feedback and correction because they distribute decision-making and impose consequences through prices, competition, profit, and loss (Hayek, 1945).

The issue of soft budget constraints deserves special attention. Kornai used this concept to describe a situation in which enterprises expect to be rescued when they fail, making losses less disciplining (Kornai, 1980, 1992). Soft budget constraints destroy the informational meaning of loss. If a firm loses money but expects subsidies, credit, or political protection, then losses no longer signal that resources should be reallocated. In a planned economy, soft budget constraints are common because enterprises are politically embedded and employment, output, and regional stability matter to authorities. AI cannot solve this if the political commitment to let inefficient units fail is absent. The algorithm may identify inefficiency, but officials may override it. Or worse, the algorithm may be designed to preserve politically favored firms. Calculation requires not only information but institutional willingness to act on information.

This point applies equally to modern industrial policy and state capitalism. A government may use AI to identify strategic sectors, allocate subsidies, and monitor performance. Some of this may be useful. But if politically connected firms learn how to present themselves as

strategic, innovative, green, resilient, or nationally important, they may capture the system. AI may supply technical legitimacy to discretionary allocation. The planner can point to models, forecasts, and scores, but the underlying process may still be shaped by lobbying, rent-seeking, and political favoritism. Tullock's and Buchanan's public choice insights are relevant here: government decision-makers are not disinterested social welfare maximizers; they operate under incentives, pressures, and institutional constraints (Buchanan and Tullock, 1962; Tullock, 1965). AI does not make politics disappear. It may make political allocation appear more scientific.

The same danger applies to data collection itself. Actors may influence what data are collected, how variables are defined, how models are trained, and how results are interpreted. In a comprehensive AI-planning system, the struggle for resources becomes a struggle to shape the algorithm. Industries, regions, ministries, unions, firms, and interest groups will seek favorable classifications and weights. Should an industry be classified as essential? Should a region be classified as disadvantaged? Should a technology be classified as green? Should a firm be classified as systemically important? These are not purely technical choices. They determine access to capital, labor, permissions, and subsidies. The politics of planning becomes the politics of model design.

This is one reason why transparency is not a complete solution. One might argue that AI planning can be made transparent, auditable, and accountable. These are desirable goals. But transparency does not remove strategic behavior. If the algorithm is transparent, actors can game it more easily. If it is opaque, actors cannot contest it effectively. If it changes constantly, it becomes unpredictable and may undermine planning. If it is stable, it becomes exploitable. This is a general problem of algorithmic governance. Markets also face manipulation, but no single algorithm determines all allocation. In a centralized system, the allocation rule becomes a prize to be captured or gamed.

Another problem is that AI systems may reward legibility over value. Local actors may devote resources to making their activities visible to the algorithm rather than genuinely productive. A teacher may document measurable outputs rather than cultivate understanding. A doctor may optimize records rather than patient trust. A firm may improve compliance dashboards rather than quality. A researcher may produce measurable publications rather than deep discoveries. A city may generate smart-city data rather than livability. The old bureaucratic pathology of paperwork becomes the new algorithmic pathology of datawork. The system becomes increasingly sophisticated at measuring what actors increasingly perform for measurement.

This is closely related to James C. Scott's analysis of legibility in state projects. Scott argues that high-modernist states often simplify complex social realities in order to administer them, but these simplifications can destroy local knowledge and practical order (Scott, 1998). Although Scott's framework differs from Hayek's, the affinity is clear. AI planning may be the most advanced form of legibility-seeking governance. It can see more than previous states, but what it sees is still structured by administrative categories. If local actors reshape themselves to fit those categories, the society becomes more legible but not necessarily more

knowledgeable, productive, or free.

The problem of truthful preference revelation is especially acute. A consumer's preference cannot be inferred simply from statements or even from observed behavior if the institutional environment is distorted. If goods are rationed, people overstate need. If goods are subsidized, consumption rises artificially. If prices are absent or controlled, demand loses its normal meaning. If behavior is surveilled, people may conceal politically sensitive preferences. If recommendations shape choices, observed demand may reflect algorithmic nudging rather than autonomous desire. Thus an AI system may observe vast quantities of behavior while misunderstanding the underlying valuations. Hayek's argument for prices is partly that prices reveal relative valuations under scarcity. Without scarcity-facing choice, preference data become unreliable (Hayek, 1945).

This also applies to labor preferences. Workers may report willingness to perform certain jobs if compensation, risk, dignity, location, family obligations, and future opportunity are not fully considered. An AI system may infer that a worker is suited to a task, but the worker may possess private knowledge about aspirations, fears, relationships, health, or moral commitments. If labor allocation becomes algorithmic, workers may learn to manipulate profiles, hide abilities, exaggerate constraints, or conform to categories. A free labor market is not perfect, but it allows individuals to act on private knowledge without fully disclosing it. They can quit, move, retrain, bargain, and choose. These actions reveal information in ways that administrative assignment cannot easily replicate.

The incentive problem also affects local officials and planners themselves. A central AI system does not operate in a vacuum. It is built, maintained, interpreted, and governed by human beings and institutions. Officials may suppress unfavorable outputs, adjust objectives, reinterpret results, or selectively enforce recommendations. Model builders may embed assumptions. Agencies may compete for jurisdiction. Politicians may demand exceptions. Bureaucrats may prefer rules that protect themselves from blame. Thus "the AI planner" is not a single neutral mind. It is an institutional apparatus populated by actors with incentives. Hayek's critique of central planning applies not only to the planner's knowledge but also to the planner's political situation (Hayek, 1944, 1945).

In a democratic society, AI planning faces additional tensions. If the algorithm makes allocation decisions, citizens will demand accountability. But democratic accountability may conflict with algorithmic optimization. Voters may demand subsidies, protections, jobs, or regional investments that the model judges inefficient. Politicians may override the model. Alternatively, officials may use the model to avoid responsibility: "the algorithm decided." In either case, AI does not remove politics. It changes the form of political conflict. In an authoritarian society, the problem is even more severe: AI planning may become fused with surveillance, repression, and loyalty scoring. In both cases, the knowledge problem becomes entangled with the problem of power.

A defender of AI planning may reply that markets also generate strategic manipulation. Advertising manipulates consumers. Firms hide defects. Financial markets produce bubbles. Workers may shirk. Managers may pursue private benefits. Monopolies may restrict output.

These criticisms are valid. Hayekian economics should not romanticize markets. But the relevant question is comparative. Markets contain decentralized forms of correction: competition, exit, entry, litigation, reputation, bankruptcy, short selling, journalism, regulation, and consumer choice. No mechanism is perfect, but errors are not necessarily unified under one plan. In comprehensive AI planning, the central system's errors and incentives can become systemic. When the algorithm mismeasures, many actors adjust to the mismeasurement at once.

This is why polycentricity matters. Elinor Ostrom's work, though developed in a different context, is helpful here. Polycentric systems contain multiple centers of decision-making, allowing experimentation, local adaptation, mutual monitoring, and correction across levels (Ostrom, 1990, 2010). A Hayekian AI economy would be polycentric rather than monocentric. It would allow many actors to use AI tools, generate knowledge, test plans, and correct errors. A centralized AI-planning system, by contrast, risks turning one model into a society-wide bottleneck. The incentive to manipulate that bottleneck becomes enormous.

The conclusion of this section is therefore straightforward. Artificial intelligence improves monitoring, prediction, and data analysis, but it does not eliminate the incentive problem. Knowledge must be elicited from agents who have interests, fears, ambitions, and strategies. The more allocation depends on measured indicators, the more people will adapt to those indicators. The more comprehensive the algorithm, the more valuable it becomes to manipulate the algorithm. The more centralized the system, the more systemic its errors become. The socialist planning problem was never merely that planners could not see enough; it was that what they saw was shaped by distorted incentives. AI can see more, but it cannot make strategic behavior disappear.

Hayek's price system remains powerful because it elicits knowledge through decentralized, costly, and consequence-bearing action. Mises's calculation argument remains powerful because genuine prices arise from property and exchange, not from declarations or reports. Kornai's shortage economy remains relevant because administrative allocation creates incentives to hoard, bargain, and distort. Public choice remains relevant because planners and politicians have incentives of their own. AI alters the technology of planning, but not the basic institutional problem: information is produced by people acting within rules. Change the rules, and the information changes. An AI-planned economy would not simply collect knowledge; it would produce a new strategic environment in which knowledge itself is shaped, hidden, gamed, and politicized.

8 Surveillance, Power, and the Political Economy of AI Planning

The fifth Hayekian rebuttal concerns power. The knowledge problem is not merely epistemic; it is political. A society that attempts to solve the knowledge problem through comprehensive artificial intelligence must also decide who owns the data, who controls the models, who

defines the objectives, who has authority to override the system, who bears the costs of error, and who may exit from the resulting order. These are not technical details. They determine whether AI becomes a tool for decentralized empowerment or a machinery of centralized control. Hayek's argument in "The Use of Knowledge in Society" shows why dispersed knowledge cannot be fully collected and used by one authority (Hayek, 1945). His argument in *The Road to Serfdom* shows why the attempt to impose one comprehensive economic plan tends to concentrate political power and restrict individual freedom (Hayek, 1944). In the age of AI, these two arguments converge. The more ambitious the plan, the greater the data appetite; the greater the data appetite, the stronger the tendency toward surveillance; the stronger the surveillance apparatus, the more dangerous the concentration of power.

The pro-AI planner might say that artificial intelligence need not be coercive. It could simply collect information, infer preferences, forecast needs, and allocate resources more efficiently. But comprehensive planning requires much more than voluntary surveys or aggregate statistics. If the planner must know what people want, what they can produce, how they might respond to incentives, which resources they possess, which skills they have, where bottlenecks will arise, and how preferences are changing, then the planner must observe economic life at great depth. It must monitor consumption, production, movement, communication, work, saving, investment, health, education, and perhaps even social relationships. In the older planning systems, this information was partially collected through reports, ministries, local officials, production targets, and rationing systems. In an AI-planned economy, the same impulse would be technologically upgraded through digital payments, sensors, cameras, platform records, biometric systems, location tracking, workplace monitoring, and predictive analytics.

This creates a fundamental paradox. The planner seeks knowledge in order to coordinate society, but the act of extracting knowledge changes the society being coordinated. People behave differently when they are watched. They conceal, conform, perform, and strategically adapt. Surveillance may produce more data, but not necessarily more truthful knowledge. It may produce a society that is more legible to the center and less capable of genuine creativity, trust, dissent, and experimentation. This is directly connected to Hayek's core insight. The knowledge that matters in a free society is not merely stored in persons; it is expressed through action under conditions of liberty. If people cannot act freely, the knowledge generated by their actions changes. The attempt to solve the knowledge problem through surveillance may therefore destroy part of the very knowledge it seeks to obtain.

Hayek's warning in *The Road to Serfdom* remains relevant precisely because central planning cannot avoid ranking ends. A comprehensive economic plan must decide not only how to allocate means, but which ends deserve priority. Should more resources go to housing, health, defense, climate transition, consumption, research, leisure, equality, religious life, family support, national prestige, or artistic excellence? In a market order, many such choices are made by individuals and associations within general rules. In a planned order, they become matters of collective decision and administrative enforcement.

Hayek argued that planning tends toward coercion because individuals do not agree on one hierarchy of ends; to execute a unified plan, authorities must impose priorities on those who dissent (Hayek, 1944). AI does not remove this problem. It may hide it behind optimization language, but every optimization system requires an objective function.

The objective-function problem is perhaps the clearest way to see why AI planning is political. An algorithm does not allocate resources in the abstract. It maximizes or balances something: output, welfare, equality, health, security, carbon reduction, employment, innovation, social stability, national power, or some weighted combination. The choice of weights is not dictated by data. It is a moral and political decision. Even if the algorithm is technically perfect, it cannot tell society whether growth should be sacrificed for equality, whether privacy should be sacrificed for health, whether consumption should be sacrificed for climate goals, whether local communities should be sacrificed for national efficiency, or whether individual liberty should be sacrificed for social stability. These questions precede computation. They are not solved by intelligence, artificial or otherwise.

This is why technocratic planning often disguises politics as expertise. A planning authority may say that the model has identified the optimal allocation. But “optimal” always means optimal relative to specified objectives, constraints, and value judgments. In a pluralistic society, these value judgments are contested. A person may prefer less income and more autonomy. A community may prefer local control to efficiency. A religious group may reject goods that others value. An artist may prefer poverty with freedom to comfort under supervision. An entrepreneur may prefer the right to fail to the security of assignment. A central AI system can classify these preferences, but it cannot legitimate the hierarchy among them. Hayek’s liberalism rests on the recognition that no single scale of value can organize the whole of social life without suppressing legitimate diversity (Hayek, 1944, 1960, 1973).

This connects the AI-planning debate to the broader problem of liberal freedom. Markets are not merely efficient because they process information; they are politically valuable because they reduce the need for centralized agreement about ends. They allow people to pursue many plans simultaneously. Individuals can buy, sell, work, save, donate, invest, experiment, fail, move, and associate without needing approval from a central authority for every major choice. Prices coordinate these plans without requiring moral unanimity. Of course, markets require law, regulation, property rights, public goods, and limits on coercion, fraud, and external harm. But they preserve a wide sphere of decentralized choice. Comprehensive AI planning would shrink that sphere by making more choices objects of administrative optimization.

The danger becomes even greater when data and allocation are joined. If the same authority that observes behavior also allocates resources, then knowledge becomes power in the most direct sense. A state that knows a person’s consumption, travel, communications, political views, health, employment, education, debts, social network, and behavioral tendencies can allocate opportunities accordingly. It can reward compliance, penalize dissent, restrict mobility, shape information exposure, and condition access to goods. Such

a system need not use crude repression. It can govern through nudges, scores, permissions, rankings, delays, and differential access. AI planning therefore may produce a softer but deeper form of control: not merely commands, but continuous behavioral steering.

Shoshana Zuboff's concept of surveillance capitalism is useful here, even though her primary target is private platform power rather than socialist planning. Zuboff argues that digital firms extract behavioral data and use it to predict and shape future behavior for profit (Zuboff, 2019). The relevance to AI planning is clear. If private platforms can use data extraction to influence consumption, attention, and social behavior, a state or central planning authority with broader powers could do the same across the whole economy. The problem is not only that the planner might know too much; it is that the planner might use knowledge to shape what people come to want. At that point, planning no longer responds to preferences. It manufactures them.

This is a fundamental challenge to the claim that AI can infer preferences and allocate accordingly. If the system both observes and shapes preferences, then revealed behavior cannot be treated as independent evidence of welfare. A recommendation algorithm may infer that users prefer certain content, but if the algorithm itself trained users to prefer that content, the inference is circular. A planning system may infer that citizens prefer stability, convenience, or officially approved goods, but if dissenting alternatives are hidden, taxed, delayed, or discouraged, the observed preference is partly produced by the plan. Hayek's defense of spontaneous order is partly a defense against this circularity. In a free order, preferences and plans emerge from many sources. In a centralized order, they risk becoming products of the system that claims merely to satisfy them.

The authoritarian potential of AI planning is not hypothetical. Modern states already use digital technologies for surveillance, censorship, policing, welfare administration, border control, tax enforcement, and social monitoring. Some uses are legitimate; others are dangerous. The point is not that all state data systems are oppressive. The point is that comprehensive economic planning would require a scale and integration of data that dramatically increases the temptation and capacity for political control. Hayek's fear was that economic control becomes political control because whoever controls the means by which people live also controls much of life itself (Hayek, 1944). AI intensifies this danger because it permits more granular and predictive forms of control than earlier planning systems could achieve.

This risk also applies to private power. AI-enabled centralization may occur not only through the state but also through dominant corporations and platforms. Large digital firms already mediate access to markets, information, cloud infrastructure, app stores, advertising, logistics, and social communication. They can rank, recommend, exclude, promote, demote, price-discriminate, and collect behavioral data at scale. Brynjolfsson and Hitzig's argument that AI may increase firm size and industry concentration therefore has political implications as well as economic ones (Brynjolfsson and Hitzig, 2025). If AI shifts decision-making upward inside firms and strengthens central control by headquarters or platforms, then Hayekian concerns about dispersed knowledge must be extended to corporate hierarchy.

The problem is not only “the state versus the market,” but centralized knowledge power versus polycentric social order.

A large platform may not be a socialist planning board, but it can become a private planning authority over a domain of economic life. It can decide which sellers are visible, which products are recommended, which workers receive tasks, which users are trusted, which speech circulates, and which fees apply. It can use AI to optimize the ecosystem according to its own objective function. This may produce efficiency, convenience, and innovation. But it also raises questions of entry, dependency, contestability, and voice. If sellers, workers, or consumers cannot easily exit, then platform governance begins to resemble administrative control. Thus the AI age requires a Hayekian analysis of both public and private centralization.

This point is important because a naïve pro-market response might celebrate AI-driven corporate centralization simply because it occurs within capitalism. But Hayek’s deepest argument concerns the use of dispersed knowledge, not merely the formal distinction between public and private ownership. If a few platforms become bottlenecks through which vast amounts of economic knowledge must pass, then the epistemic advantages of decentralization may be weakened. Competition, entry, interoperability, data portability, antitrust policy, open standards, and decentralized AI tools therefore become central to a modern Hayekian institutional agenda. The goal is not to prevent organization, but to prevent any single organization from monopolizing the conditions of discovery.

The political economy of AI planning also raises the problem of error. Centralized systems can make large-scale mistakes. If a decentralized market contains many firms and investors, errors are distributed. Some firms overinvest while others underinvest. Some entrepreneurs misjudge demand while others discover opportunities. Some localities fail while others succeed. The system is not perfect, but mistakes are partially compartmentalized. A comprehensive AI system, by contrast, may propagate one error across the whole economy. If the model misestimates energy demand, undervalues a technology, overprioritizes a sector, suppresses a product, misclassifies risk, or rewards the wrong indicator, the consequences become systemic. The concentration of intelligence also concentrates error.

This is not merely a technical problem of model accuracy. It is a political problem of accountability. Who is responsible when the AI plan fails? The engineers who built the model? The officials who set the objectives? The politicians who approved it? The agencies that implemented it? The data providers? The algorithm itself? In markets, responsibility is imperfect but often more traceable through ownership, contract, liability, profit, and loss. In bureaucratic systems, responsibility can become diffused. In algorithmic systems, it may become even more obscure. The phrase “the model recommended it” can become a shield against accountability. Hayek’s critique of planning includes precisely this concern: when decisions are centralized, individual responsibility may weaken even as power increases (Hayek, 1944).

A further danger is that AI planning may undermine the rule of law. Hayek’s political theory places great weight on general, abstract, and predictable rules rather than

discretionary commands (Hayek, 1960, 1973). Markets require a legal framework in which individuals can form plans. Comprehensive AI governance, however, may favor continuous discretionary adjustment. If the algorithm can personalize prices, permissions, taxes, subsidies, regulations, credit terms, and access conditions, then general rules may give way to individualized administration. Such personalization may appear efficient because it tailors policy to circumstances. But it may also undermine equality before the law. Citizens may no longer know the rules that govern them. They may be treated differently according to opaque risk scores or behavioral predictions. The result is not spontaneous order under general rules, but managed order under continuous administrative discretion.

This is especially dangerous when predictive systems are used to allocate opportunities. If AI predicts that a person is unlikely to repay a loan, unlikely to succeed in education, likely to commit a crime, likely to be unhealthy, or likely to be politically unreliable, institutions may restrict access before the person acts. Prediction then becomes preemptive governance. The individual is treated according to a statistical profile rather than as a responsible agent. Such systems may improve efficiency in some contexts, but if generalized across society they threaten the liberal idea of personhood. Hayek's order depends on individuals being free to use their own knowledge and revise their own plans. Predictive administration risks freezing people into categories generated by past data.

The pro-AI planner may respond that democratic oversight can prevent abuse. Democratic control is essential, but it is not sufficient by itself. Democratic majorities may also desire planning, surveillance, redistribution, censorship, industrial favoritism, or behavioral control. Interest groups may capture the system. Voters may not understand complex algorithms. Agencies may hide behind technical expertise. Moreover, democracy works best when the sphere of collective decision is limited enough that disagreement does not become total conflict. Hayek's liberal constitutionalism seeks to preserve democracy by limiting the range of discretionary power subject to political struggle (Hayek, 1960, 1973). Comprehensive AI planning moves in the opposite direction: it expands the domain of politically controlled allocation.

Public choice theory reinforces this concern. Buchanan and Tullock argued that political actors respond to incentives and that collective decision-making can generate rent-seeking, logrolling, and concentrated benefits with dispersed costs (Buchanan and Tullock, 1962). Tullock's work on rent-seeking showed that political allocation creates incentives for groups to expend resources capturing favors rather than producing value (Tullock, 1965). AI does not eliminate these incentives. It may provide new tools for targeting benefits, justifying subsidies, monitoring constituencies, and rewarding supporters. Industrial policy guided by AI may still be captured by firms that shape data, lobby for favorable model assumptions, or present themselves as strategically essential. Algorithmic planning can become a new language for old politics.

This is why the ownership and governance of AI systems matter. If AI is monopolized by the state, planning risks authoritarianism. If AI is monopolized by a few corporations, markets risk private domination. If AI is distributed widely, it may strengthen decentralized

knowledge use. A Hayekian approach should therefore favor institutional arrangements that prevent the concentration of data and decision power. This may include competition policy, open-source AI, interoperable systems, data portability, privacy protections, decentralized computation, strong property and contract rights, and legal limits on state surveillance. The purpose is not to stop AI but to prevent the emergence of a single center of social cognition.

The contrast between monocentric and polycentric order is crucial. Elinor Ostrom's work on polycentric governance shows that complex societies often solve problems through multiple overlapping centers of decision-making rather than through one hierarchy (Ostrom, 1990, 2010). Although Ostrom's research focused especially on common-pool resources and institutional diversity, her framework complements Hayek's. Polycentric systems allow local knowledge, experimentation, mutual monitoring, and institutional learning. AI could be used polycentrically: local governments, firms, households, cooperatives, researchers, courts, regulators, and entrepreneurs could use AI tools to improve decisions while remaining subject to feedback and contestation. The danger arises when AI is used monocentrically, as the intelligence of one comprehensive plan.

The normative conclusion follows. AI planning is not merely a question of whether machines can calculate better than bureaucrats. It is a question of what kind of society we want. A society can use AI to help persons make decisions, or it can use AI to make persons more governable. It can distribute intelligence, or it can centralize it. It can preserve privacy and pluralism, or it can make social life transparent to authority. It can encourage experimentation, or it can optimize conformity. It can strengthen markets by improving search and coordination, or it can replace markets with algorithmic allocation. These choices are institutional and political.

Hayek's continuing relevance lies in his insistence that social order cannot be reduced to conscious control. The more complex society becomes, the more dangerous it is to assume that a central intelligence can see, rank, and direct all relevant activity. Artificial intelligence increases the power of conscious control, but it does not eliminate the limits of such control. Indeed, by making control more tempting, AI makes Hayek's warning more urgent. The dream of an omniscient planner is no longer a fantasy of bureaucratic rationalism alone; it is becoming a technological temptation.

The conclusion of this section is therefore clear. Comprehensive AI planning would require extensive data extraction, centralized objective-setting, administrative discretion, and political authority. These requirements create risks of surveillance, preference manipulation, rent-seeking, soft coercion, systemic error, and loss of pluralism. AI can and should be used in many domains, including public administration and firm organization. But when AI is joined to comprehensive economic control, the knowledge problem becomes a power problem. The planner who claims to know society may also acquire the power to shape society. Hayek's defense of the price system and spontaneous order is therefore inseparable from a defense of liberal limits on power. In the age of artificial intelligence, the old question "who plans?" becomes inseparable from the new question "who sees?"

9 Firms as Islands of Planning Within Markets: AI, Hierarchy, Platforms, and the Coasean Boundary

A serious discussion of AI and Hayek must confront an apparent paradox. If Hayek is right that decentralized markets make better use of dispersed knowledge than central planning, why do firms exist at all? A firm is, in an important sense, a planned order. Within a firm, many resource allocations are not made through market exchange but through managerial direction, budgets, routines, contracts, and hierarchy. Workers usually do not auction every task internally. Departments often receive allocations through administrative decisions. Managers plan production, marketing, logistics, hiring, inventory, pricing, and investment. Modern corporations may coordinate millions of transactions, thousands of workers, and complex global supply chains through centralized information systems. If planning is inherently inferior, the existence of large firms appears to be a problem for Hayekian theory.

The answer begins with Coase. In “The Nature of the Firm,” Coase argued that firms exist because using the price mechanism is itself costly. Market exchange requires searching for trading partners, negotiating contracts, enforcing agreements, adapting to contingencies, and coordinating repeated activities. When the costs of market contracting exceed the costs of internal organization, activities may be brought inside the firm (Coase, 1937). Williamson later developed this insight into a broader theory of transaction-cost economics, arguing that markets and hierarchies are alternative governance structures suited to different circumstances, especially where asset specificity, uncertainty, frequency of exchange, and opportunism matter (Williamson, 1975, 1985). Firms are therefore not refutations of the market order. They are institutions that economize on certain transaction costs while remaining embedded in a larger market environment.

This Coasean framework is essential for understanding artificial intelligence. AI may reduce some costs of internal organization. It can monitor operations, forecast demand, compare performance, automate routine decisions, coordinate logistics, detect fraud, schedule labor, manage inventories, and integrate dispersed information across many units. These capacities may enlarge the efficient scale of the firm. Brynjolfsson and Hitzig argue precisely that AI can codify some local knowledge, reduce the informational advantage of local managers, shift decision-making upward, and thereby increase firm size and industry concentration in some settings (Brynjolfsson and Hitzig, 2025). This is a real and important modification of the Hayekian landscape. AI may make hierarchy more capable. It may make headquarters smarter. It may permit large organizations to coordinate more effectively than before.

But this does not imply that the whole economy can or should be organized as one firm. The reason is that firms use planning within markets, not instead of markets. A corporation can internally plan because it relies on external prices for labor, land, capital, raw materials, energy, components, transportation, credit, and consumer goods. Its managers can calculate

costs because suppliers quote prices. They can evaluate revenues because consumers buy or refuse to buy. They can assess investment because capital markets impose costs and expectations. They can judge failure because losses appear. They can compare alternatives because the market outside the firm generates opportunity-cost signals. Hayek's price system surrounds and disciplines the firm even when the firm uses hierarchy internally (Hayek, 1945).

This is why the analogy between a large firm and a socialist economy is misleading. A firm is an island of planning in an ocean of market prices. A socialist economy attempts to turn the whole ocean into one island. Once the entire economy is internalized into a single hierarchy, the external price environment that made internal planning possible disappears. The firm can plan because it does not have to generate all prices internally. It can use market prices as inputs into its own calculation. If one imagines a whole society as one giant firm, the problem becomes radically different. There is no external market left to discipline the plan. This was the core Mises-Hayek point: internal planning may work locally because it borrows information from the wider market order, but comprehensive planning cannot borrow from a market it has abolished (Mises, 1920; Hayek, 1945).

Artificial intelligence does not change this institutional logic. Amazon's logistics algorithms, Walmart's inventory systems, Toyota's production methods, Apple's supply-chain management, and cloud-computing platforms are impressive examples of large-scale planning. But their success depends on markets. Amazon observes actual consumer purchases, pays market wages, contracts with suppliers, uses capital markets, competes with other retailers, and faces consumer exit. Walmart's inventory algorithms depend on supplier prices, consumer demand, transport costs, and competitive pressure. Apple's supply-chain coordination relies on global markets for components, labor, intellectual property, and consumer products. These firms are not evidence that markets are obsolete. They are evidence that planning can be powerful when disciplined by market-generated information.

The AI age may therefore shift the Coasean boundary without eliminating the Coasean problem. Some tasks formerly governed by market contracting may move inside firms because AI lowers internal coordination costs. Other tasks may move outside firms because AI lowers search, contracting, and verification costs in markets. The direction is not predetermined. AI may enable larger firms in sectors where data scale, network effects, and coordination complementarities dominate. But AI may also empower small firms, freelancers, entrepreneurs, and local producers by giving them access to tools once available only to large organizations. A small business can now use AI for accounting, marketing, translation, legal drafting, customer service, design, forecasting, and software development. Thus AI may centralize some activities and decentralize others. The institutional outcome depends on competition, regulation, data access, interoperability, and entrepreneurial entry.

Platforms complicate the picture further. Digital platforms are neither ordinary firms nor ordinary markets. They are governance structures that host market-like interactions while controlling rules, visibility, pricing systems, reputation mechanisms, search rankings, access conditions, and data flows. Ride-sharing platforms coordinate drivers and riders through

algorithmic matching and dynamic pricing. E-commerce platforms coordinate buyers, sellers, advertisers, warehouses, payment systems, and logistics. App stores coordinate developers and users while enforcing standards and fees. Cloud platforms coordinate computing resources across firms. These systems show that AI does not simply strengthen hierarchy; it creates hybrid forms of algorithmic governance.

From a Hayekian perspective, platforms are ambivalent. On one side, they can improve coordination by reducing search costs, creating trust mechanisms, aggregating dispersed supply and demand, and enabling small sellers to reach broad markets. They can make markets thicker and more efficient. On the other side, they can become bottlenecks that centralize knowledge and control. A dominant platform may decide which products are visible, which sellers survive, which fees apply, which workers receive tasks, and which users are trusted. It may use AI not merely to facilitate exchange but to shape behavior according to its own objective function. Thus platform capitalism confirms the relevance of Hayek's knowledge problem while extending it. The question is no longer only whether the state should plan the economy, but whether private digital authorities can accumulate planning-like power over large domains of social life.

This point is especially important for competition policy. If AI creates strong economies of scale in data, computing, and model performance, it may increase concentration. Large firms may improve their models because they have more users; more users generate more data; better models attract still more users. This feedback loop may produce durable advantages. If dominant firms control data infrastructures, cloud platforms, model access, logistics networks, and consumer interfaces, decentralized market discovery may be weakened. New entrants may find it difficult to compete not because they lack ideas, but because they lack access to data, distribution, or computational infrastructure. Hayekian political economy therefore cannot simply defend existing large firms as "market outcomes." It must ask whether market processes remain open, contestable, and discovery-enhancing.

The proper Hayekian response to AI-enabled firm centralization is not anti-firm romanticism. Large firms can be extraordinarily productive. They can mobilize capital, coordinate research, scale innovations, standardize quality, manage supply chains, and bring complex goods to consumers. Schumpeter himself emphasized that large firms may play a major role in innovation, especially where research and development require scale (Schumpeter, 1942). The issue is whether large firms remain subject to competitive discipline and whether outsiders can challenge them. A large firm embedded in open markets is different from a protected monopoly or a platform gatekeeper insulated from entry. The first may enhance coordination; the second may suppress discovery.

This distinction also matters for internal AI governance. If headquarters uses AI to reduce waste, improve logistics, and transmit useful knowledge across units, internal centralization may be efficient. But if headquarters uses AI to suppress local judgment, impose uniform metrics, and ignore context, it may destroy valuable knowledge within the firm. The Hayekian argument applies inside organizations as well as between market and state. Local managers, workers, engineers, salespeople, and customers often possess knowledge that

headquarters cannot fully understand. AI may help transmit some of this knowledge upward, but it may also encourage overcentralization by making headquarters believe it sees more than it actually does. The danger is not planning itself, but planning that becomes blind to the knowledge it cannot measure.

Successful firms often combine central coordination with decentralized experimentation. They use headquarters for strategy, capital allocation, standards, and shared infrastructure, but they also allow local units, product teams, engineers, and managers to experiment. They run pilot projects, A/B tests, internal venture programs, regional adaptations, and customer-feedback loops. This is a firm-level version of polycentric order. AI can strengthen such systems if used to support experimentation and learning. It can weaken them if used to impose rigid algorithmic control. The best firms will not be those that centralize every decision, but those that know which decisions should be centralized and which should remain close to local knowledge.

The same lesson applies to public policy. The fact that AI improves planning inside firms does not justify comprehensive state planning. It does suggest that governments may use AI for limited and well-specified tasks: infrastructure maintenance, emergency logistics, tax fraud detection, public-health monitoring, environmental measurement, or traffic management. These are domains where objectives are relatively concrete and feedback may be measurable. But the more the state attempts to direct investment, production, labor allocation, consumption, and innovation across the whole economy, the more it faces the full Mises-Hayek problem. Limited planning inside a market order is not the same as replacing the market order.

The conclusion of this section is therefore balanced. AI may indeed enlarge hierarchy. It may increase the efficient scale of some firms, strengthen headquarters, reduce some forms of local autonomy, and create powerful platform-governance systems. Brynjolfsson and Hitzig are right that this development puts pressure on simple Hayekian claims about decentralization (Brynjolfsson and Hitzig, 2025). But the pressure is not fatal. Firms and platforms operate within, depend upon, and are disciplined by broader market institutions. Their internal planning succeeds because external prices, competition, consumer choice, labor mobility, and capital markets generate signals that no internal algorithm can fully replace. A firm can be an island of planning because it is not the whole sea.

The AI economy should therefore be judged not by whether it contains more planning or less planning in the abstract, but by whether its planning remains embedded in a competitive, open, price-guided, and polycentric order. AI used by many firms, entrepreneurs, consumers, and public institutions can strengthen decentralized coordination. AI monopolized by a few platforms or by the state can weaken it. Hayek's lesson is not that hierarchy has no place. It is that hierarchy must be limited, disciplined, and surrounded by institutions that preserve dispersed knowledge, rivalry, exit, and discovery. In the age of artificial intelligence, the essential question is not whether machines can help organizations plan. They can. The essential question is whether the society that uses them remains open enough for knowledge to be discovered outside the plan.

10 Conclusion: Hayek After Artificial Intelligence

The rise of artificial intelligence forces Hayekian political economy to become more precise. It is no longer enough to say, in a simple way, that central planners cannot process enough information. In many domains, AI systems can process vastly more information than local human decision-makers. They can aggregate data from millions of transactions, detect patterns invisible to ordinary managers, forecast demand, optimize logistics, monitor supply chains, automate routine decisions, and even codify some forms of previously tacit knowledge. AI can strengthen corporate headquarters, enlarge the feasible scale of firms, improve public administration, assist scientific research, and support more sophisticated forms of industrial policy. A serious defense of Hayek must therefore begin by admitting that the technological environment has changed. Hayek wrote in 1945; we now live in a world of digital platforms, real-time data, machine learning, cloud computing, sensors, and large language models. The knowledge problem must be restated for this new world.

But restating Hayek is not the same as abandoning him. The central argument of this paper has been that artificial intelligence modifies the boundary between centralization and decentralization, but it does not abolish the knowledge problem. Hayek's deepest claim was never that planners lacked sufficiently powerful calculating machines. His claim was that the knowledge required for economic coordination is dispersed, tacit, contextual, subjective, changing, incentive-laden, and often generated only through decentralized action (Hayek, 1945). This knowledge is not simply "data" waiting to be collected. It is embedded in local circumstances, practical judgment, entrepreneurial imagination, consumer experimentation, institutional trust, and action under real constraints. AI can process recorded traces of human behavior, but recorded traces are not identical to knowledge. The visible is not the whole of the real; the measurable is not the whole of the relevant.

The first conclusion, therefore, is that data are not knowledge. Modern AI systems excel at extracting patterns from data, but economic knowledge includes interpretation, relevance, context, and judgment. A shopkeeper's understanding of customers, a worker's knowledge of a machine, an entrepreneur's sense of opportunity, a consumer's changing preference, or a manager's reading of trust and morale may not be fully captured in digital form. Polanyi's insight that "we can know more than we can tell" remains relevant even though modern AI can now learn some things from examples rather than explicit rules (Polanyi, 1966; Autor, 2014). AI weakens some older versions of the tacit-knowledge argument, but it does not eliminate tacitness altogether. Much human knowledge remains situated, embodied, relational, and inseparable from responsibility.

The second conclusion is that prediction is not discovery. AI is powerful when the future resembles the past or when enough data exist to estimate relevant patterns. But capitalism is not merely a prediction system. It is a discovery process. New goods, new technologies, new uses of resources, new organizational forms, and new preferences are not always hidden in existing data. They are often created through entrepreneurial action, competition, experimentation, and failure. Hayek's account of competition as a discovery

procedure, Kirzner's theory of entrepreneurial alertness, and Schumpeter's theory of new combinations all remain essential for understanding why decentralized markets cannot be replaced by prediction alone (Hayek, 1968; Kirzner, 1973; Schumpeter, 1911, 1942). AI can help entrepreneurs discover. It can widen search and reduce the cost of experimentation. But it cannot fully replace the open-ended social process through which new possibilities are generated and tested.

The third conclusion is that simulated prices are not market prices. Mises's calculation argument remains indispensable. A central algorithm can generate shadow prices, estimate demand, simulate allocation, and adjust administrative prices. But market prices are not merely numbers. They are institutional outcomes of exchange under private property, competition, scarcity, profit, and loss (Mises, 1920, 1949). They reveal opportunity costs because people act under constraints and bear consequences. A simulated price may be useful inside a firm, agency, or model, but it does not possess the same epistemic status as a price generated by rivalrous bidding among owners, consumers, workers, investors, and entrepreneurs. AI may improve calculation where genuine market prices already exist. It cannot by itself create the institutional conditions that make prices meaningful.

The fourth conclusion is that monitoring does not solve incentives. A central AI system may see more than earlier planners, but the information it receives is still produced by people acting under incentives. If resources, promotions, permissions, or punishments depend on algorithmic indicators, people will learn to manipulate those indicators. The history of socialist planning, as analyzed by Kornai, shows that administrative systems generate hoarding, bargaining, distorted reporting, soft budget constraints, and plan-fulfillment behavior (Kornai, 1980, 1992). AI may detect some manipulation, but it also creates new forms of metric gaming and algorithmic strategic behavior. Knowledge is not passively collected; it is elicited or distorted by institutions. Change the rules, and the information changes.

The fifth conclusion is that the knowledge problem becomes a power problem when AI is joined to comprehensive planning. To collect all relevant knowledge, a planner must observe more of social life. The desire for better planning therefore generates a desire for more data, and the desire for more data easily becomes surveillance. Hayek's warning in *The Road to Serfdom* remains urgent: comprehensive economic planning tends to concentrate political power because it requires the imposition of a hierarchy of ends on people who do not share one scale of values (Hayek, 1944). AI does not remove that problem. Every algorithmic plan requires an objective function. Every objective function embodies moral and political choices. The danger of AI planning is that these choices may be hidden behind technical language, while the power to observe, classify, predict, and allocate becomes centralized.

The sixth conclusion is that firms do not refute Hayek. Firms are indeed planned orders, and AI may enlarge the domain of hierarchy. Coase and Williamson already explained why some activities are more efficiently coordinated within firms than through market contracting (Coase, 1937; Williamson, 1975, 1985). Brynjolfsson and Hitzig are right that AI may shift the efficient boundary of the firm by codifying some local knowledge and increasing

headquarters' capacity to coordinate (Brynjolfsson and Hitzig, 2025). But firms plan within markets. They rely on external prices, consumer choice, labor mobility, capital markets, suppliers, competitors, and profit-loss discipline. A firm can be an island of planning because it is surrounded by a sea of market signals. Turning the whole economy into one firm would destroy the external price environment that makes internal planning possible.

The final conclusion is institutional. The question is not whether AI should be used. It should be used widely. The question is whether AI will be used polycentrically or monocentrically. A polycentric AI order would distribute artificial intelligence across households, firms, entrepreneurs, researchers, consumers, local governments, courts, regulators, and civil society. It would help many actors use their own knowledge better. It would strengthen search, learning, experimentation, and coordination within a market order. A monocentric AI order would concentrate data, models, and decision rights in a state agency, dominant platform, or administrative hierarchy. It would convert AI from a tool of discovery into a tool of control. The Hayekian answer is not anti-technology; it is anti-monopoly over social knowledge.

Artificial intelligence therefore makes Hayek more relevant, not less. It forces us to distinguish more carefully between data and knowledge, prediction and discovery, shadow prices and market prices, monitoring and truthful revelation, optimization and legitimacy, hierarchy and social order. The more powerful AI becomes, the more important it is to preserve institutions that allow dispersed human beings to act, experiment, dissent, fail, and discover. The danger is not that machines will think. The danger is that political and corporate authorities will imagine that because machines can process immense quantities of data, society itself can now be treated as an object of comprehensive administration.

Hayek's enduring lesson is that no mind, whether human or artificial, can substitute for the social process through which knowledge is generated, communicated, tested, and revised. The price system, competition, private property, profit and loss, entrepreneurship, and freedom of choice are not obsolete relics from a pre-digital age. They are institutional mechanisms for using knowledge that cannot be fully centralized. AI can make these mechanisms more powerful when it serves them. It can make society more brittle and less free when it tries to replace them.

The proper future is therefore not algorithmic socialism, nor a nostalgic return to pre-digital laissez-faire. It is an open, competitive, liberal, and technologically sophisticated order in which AI augments decentralized intelligence rather than absorbing it. Such an order would use AI to improve firms, empower consumers, assist entrepreneurs, strengthen scientific discovery, and help governments perform limited public tasks more effectively. But it would preserve the market as the central discovery procedure of economic life. It would remember that society's intelligence is not located in one center. It is distributed across persons, practices, institutions, prices, experiments, failures, and freedoms. That was Hayek's insight in 1945. It remains true in the age of artificial intelligence.

11 References

- Acemoglu, Daron. 2023a. “Hayek vs. AI Socialism: Would AI-enabled Centralized Control Work?” *The International Economy* 37, no. 2, Spring: 48–49.
- Acemoglu, Daron. 2023b. “Would AI-Enabled Communism Work?” Project Syndicate, June 28, 2023.
- Acemoglu, Daron, and Simon Johnson. 2023. *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. New York: PublicAffairs.
- Arrow, Kenneth J. 1985. “The Economics of Agency.” In John W. Pratt and Richard J. Zeckhauser, eds., *Principals and Agents: The Structure of Business*, 37–51. Boston: Harvard Business School Press.
- Autor, David H. 2014. “Polanyi’s Paradox and the Shape of Employment Growth.” NBER Working Paper No. 20485. Cambridge, MA: National Bureau of Economic Research.
- Boettke, Peter J. 2001. *Calculation and Coordination: Essays on Socialism and Transitional Political Economy*. London: Routledge.
- Brynjolfsson, Erik, and Zoë Hitzig. 2025. “AI’s Use of Knowledge in Society.” In *The Economics of Transformative AI*, chapter 12. Chicago: University of Chicago Press / National Bureau of Economic Research.
- Buchanan, James M., and Gordon Tullock. 1962. *The Calculus of Consent: Logical Foundations of Constitutional Democracy*. Ann Arbor: University of Michigan Press.
- Coase, Ronald H. 1937. “The Nature of the Firm.” *Economica*, New Series, 4, no. 16: 386–405.
- Coase, Ronald H. 1960. “The Problem of Social Cost.” *Journal of Law and Economics* 3: 1–44.
- Hayek, F. A., ed. 1935. *Collectivist Economic Planning: Critical Studies on the Possibilities of Socialism*. London: George Routledge & Sons.
- Hayek, F. A. 1937. “Economics and Knowledge.” *Economica*, New Series, 4, no. 13: 33–54.
- Hayek, F. A. 1940. “Socialist Calculation: The Competitive ‘Solution.’” *Economica*, New Series, 7, no. 26: 125–149.
- Hayek, F. A. 1944. *The Road to Serfdom*. London: George Routledge & Sons.
- Hayek, F. A. 1945. “The Use of Knowledge in Society.” *American Economic Review* 35, no. 4: 519–530.
- Hayek, F. A. 1948. “The Meaning of Competition.” In *Individualism and Economic Order*, 92–106. Chicago: University of Chicago Press.
- Hayek, F. A. 1948. *Individualism and Economic Order*. Chicago: University of Chicago Press.
- Hayek, F. A. 1960. *The Constitution of Liberty*. Chicago: University of Chicago Press.
- Hayek, F. A. 1967. “The Theory of Complex Phenomena.” In *Studies in Philosophy, Politics and Economics*, 22–42. Chicago: University of Chicago Press.

- Hayek, F. A. 1968. "Competition as a Discovery Procedure." Originally published in German as "Der Wettbewerb als Entdeckungsverfahren." Kiel: Institut für Weltwirtschaft. English translation in *Quarterly Journal of Austrian Economics* 5, no. 3, 2002: 9–23.
- Hayek, F. A. 1973. *Law, Legislation and Liberty, Volume 1: Rules and Order*. Chicago: University of Chicago Press.
- Hayek, F. A. 1976. *Law, Legislation and Liberty, Volume 2: The Mirage of Social Justice*. Chicago: University of Chicago Press.
- Hayek, F. A. 1979. *Law, Legislation and Liberty, Volume 3: The Political Order of a Free People*. Chicago: University of Chicago Press.
- Kirzner, Israel M. 1973. *Competition and Entrepreneurship*. Chicago: University of Chicago Press.
- Kirzner, Israel M. 1985. *Discovery and the Capitalist Process*. Chicago: University of Chicago Press.
- Knight, Frank H. 1921. *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.
- Kornai, János. 1980. *Economics of Shortage*. 2 vols. Amsterdam: North-Holland.
- Kornai, János. 1992. *The Socialist System: The Political Economy of Communism*. Princeton: Princeton University Press.
- Lachmann, Ludwig M. 1956. *Capital and Its Structure*. London: Bell & Sons.
- Lachmann, Ludwig M. 1986. *The Market as an Economic Process*. Oxford: Basil Blackwell.
- Lange, Oskar. 1936. "On the Economic Theory of Socialism: Part One." *Review of Economic Studies* 4, no. 1: 53–71.
- Lange, Oskar. 1937. "On the Economic Theory of Socialism: Part Two." *Review of Economic Studies* 4, no. 2: 123–142.
- Lange, Oskar, and Fred M. Taylor. 1938. *On the Economic Theory of Socialism*. Edited by Benjamin E. Lippincott. Minneapolis: University of Minnesota Press.
- Lavoie, Don. 1985a. *Rivalry and Central Planning: The Socialist Calculation Debate Reconsidered*. Cambridge: Cambridge University Press.
- Lavoie, Don. 1985b. *National Economic Planning: What Is Left?* Cambridge, MA: Ballinger.
- Mises, Ludwig von. 1920. "Die Wirtschaftsrechnung im sozialistischen Gemeinwesen." *Archiv für Sozialwissenschaften* 47: 86–121. English translation: "Economic Calculation in the Socialist Commonwealth." In F. A. Hayek, ed., *Collectivist Economic Planning*, 87–130. London: George Routledge & Sons, 1935.
- Mises, Ludwig von. 1922. *Die Gemeinwirtschaft: Untersuchungen über den Sozialismus*. Jena: Gustav Fischer. English translation: *Socialism: An Economic and Sociological Analysis*. London: Jonathan Cape, 1936.
- Mises, Ludwig von. 1949. *Human Action: A Treatise on Economics*. New Haven: Yale University Press.
- Nelson, Richard R., and Sidney G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press.

- Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.
- Ostrom, Elinor. 2010. "Beyond Markets and States: Polycentric Governance of Complex Economic Systems." *American Economic Review* 100, no. 3: 641–672.
- Polanyi, Michael. 1958. *Personal Knowledge: Towards a Post-Critical Philosophy*. Chicago: University of Chicago Press.
- Polanyi, Michael. 1966. *The Tacit Dimension*. Garden City, NY: Doubleday.
- Rizzo, Mario J., and Glen Whitman. 2019. *Escaping Paternalism: Rationality, Behavioral Economics, and Public Policy*. Cambridge: Cambridge University Press.
- Schulz, Armin W. 2026. "Choice All the Way Down: AI, Central Planning, and the Hayekian Counterarguments." *The Review of Austrian Economics*. Published online April 25, 2026.
- Schumpeter, Joseph A. 1911. *Theorie der wirtschaftlichen Entwicklung*. Leipzig: Duncker & Humblot. English translation: *The Theory of Economic Development*. Cambridge, MA: Harvard University Press, 1934.
- Schumpeter, Joseph A. 1942. *Capitalism, Socialism and Democracy*. New York: Harper & Brothers.
- Scott, James C. 1998. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven: Yale University Press.
- Stiglitz, Joseph E. 2000. "The Contributions of the Economics of Information to Twentieth Century Economics." *Quarterly Journal of Economics* 115, no. 4: 1441–1478.
- Tullock, Gordon. 1965. *The Politics of Bureaucracy*. Washington, DC: Public Affairs Press.
- Williamson, Oliver E. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.
- Williamson, Oliver E. 1985. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. New York: Free Press.
- Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs.