

RESEARCH ARTICLE

# The economic institutions of artificial intelligence

Sinclair Davidson

Blockchain Innovation Hub, RMIT University, Melbourne, VIC, Australia

Email: [sinclair.davidson@rmit.edu.au](mailto:sinclair.davidson@rmit.edu.au)

(Received 16 August 2023; revised 20 December 2023; accepted 20 December 2023)

## Abstract

This paper explores the role of artificial intelligence (AI) within economic institutions, focusing on bounded rationality as understood by Herbert Simon. Artificial Intelligence can do many things in the economy, such as increasing productivity, enhancing innovation, creating new sectors and jobs, and improving living standards. One of the ways that AI can disrupt the economy is by reducing the problem of bounded rationality. AI can help overcome this problem by processing large amounts of data, finding patterns and insights, and making predictions and recommendations. This insight raises the question: can AI overcome planning problems – could it be that central planning is now a viable option for economic organisation? This paper argues that AI does not make central planning viable at either the nation-state level or the firm level, simply because AI cannot resolve the knowledge problem as described by Ludwig von Mises and Friedrich Hayek.

**Keywords:** artificial intelligence; bounded rationality; knowledge problem

**JEL Classification:** D23; D80; P11

## Introduction

The term ‘Artificial Intelligence’ (AI) – introduced by John McCarthy in 1956 – is surprisingly poorly defined. It is apparently intuitively obvious to what is meant by the term – yet there are many different definitions. In their textbook coverage of AI, Russell and Norvig (2016) are able to point to eight different definitions of what AI could be. Melanie Mitchell (2019: 20, emphasis original) summarises the position as follows:

On the scientific side, AI researchers are investigating the mechanisms of ‘natural’ (that is, biological) intelligence by trying to embed it in computers. On the practical side, AI proponents simply want to create computer programs that perform tasks as well as or better than humans, without worrying about whether these programs are actually *thinking* in the way humans think.

For our purposes, it is this latter component of Mitchell’s definition that will become important; ‘whether these programs are actually *thinking* in the way humans think’. In particular, can AI be entrepreneurial? The argument in this paper is that in order to outperform humans in some tasks, these programs will actually need to think in the way humans think. The challenge for AI researchers is that, as yet, neither they (or economists for that matter) fully understand how it is that human think.

At its core, AI operates on the principle of learning from (historical) input data to recognise patterns, make predictions, comprehend linguistic structures, and conduct image recognition tasks. This is achieved through the use of machine learning algorithms, a subset of AI, which can ‘learn’ and ‘adapt’ based on the data they process (Russell and Norvig, 2016).

AI can be categorised into two types: narrow AI, which is designed to perform a specific task and general AI, which can perform any intellectual task that a human being can do (see Mitchell, 2019 and Russell and Norvig, 2016). Currently, all existing AI technologies, including Large Language Models (LLMs) are examples of narrow AI.

General AI, also known as strong AI or Artificial General Intelligence (AGI), is a type of AI that possesses the capacity to understand, learn, adapt, and implement knowledge across a broad range of tasks, at a level equal to or indistinguishable from that of a human. It denotes a machine's ability to independently solve problems, make decisions, plan for the future, understand complex ideas, learn from experience, and apply knowledge to different domains (Russell and Norvig, 2016).

Currently, AGI remains theoretical and is mostly constrained to science fiction. HAL9000, from *2001: A Space Odyssey*, is an example of AGI. As is SkyNet from the *Terminator* series. Not all AGIs in science fiction are villains – Data in the *Star Trek* series is often portrayed as a force for moral improvement. The 'Minds' in Iain M Banks' culture series pursue the best interests of their civilisation but can also be morally ambiguous. Nonetheless, it is very likely the case that most AGIs in popular culture are villains pursuing the destruction of humanity. See Allen *et al.* (2020) for a discussion of the dangers AI may pose to humans.

While it is a long-term goal for many AI researchers, the creation of a system that exhibits human-like understanding and cognition across a diverse range of tasks is a formidable challenge and the subject of ongoing research. Mitchell (2019 – see especially pages 272–275) provides compelling arguments for the view that the development of AGI will not occur soon (if at all). In summary, her argument is that AI can (only) learn particular skills, but humans can learn to think and then apply that thinking (knowledge) to new and different situations.

Many AI technologies are already ubiquitous across the economy. In the field of finance, AI is being used to analyse historical and real-time data to detect fraud, manage risk, and trade on the basis of pricing anomalies. In supply chain management, AI can optimise operations by predicting demand, managing inventory, and streamlining logistics. In healthcare, AI is employed in medical imaging to analyse X-rays, MRIs, and CT scans, aiding doctors in diagnosing diseases and identifying abnormalities. AI is used by ride-sharing platforms to match drivers and passengers efficiently, considering factors such as distance, time, and traffic conditions, resulting in improved ride experiences and reduced wait times.

The important point is this: AI researchers have been able to create computer programs that perform some tasks better than humans are able to perform those tasks. In many instances, those computer programs are more accurate and faster than are humans. On the other hand, however, there are many tasks that AI researchers have been (to-date) unable to replicate. As Mitchell (2019: 33–34) explains:

Marvin Minsky pointed out that in fact AI research had uncovered a paradox: 'Easy things are hard.' The original goals of AI – computers that could converse with us in natural language, describe what they saw through their camera eyes, learn new concepts after seeing only a few examples – are things that young children can easily do, but, surprisingly, these 'easy things' have turned out to be harder for AI to achieve than diagnosing complex diseases, beating human champions at chess and Go, and solving complex algebraic problems. As Minsky went on, 'In general, we're least aware of what our minds do best.'

This paper explores the role of AI, with a focus on bounded rationality as proposed by Herbert Simon, and expanded by Gerd Gigerenzer and Vernon Smith, in resolving the so-called economic problem as articulated by Friedrich Hayek (1945). What impact, if any, will AI have on planning within the economy? If the problems of bounded rationality are resolved by AI, does that mean that planning is a viable mechanism to allocate resources in the economy? Is central planning now viable? What about planning within organisations? Will we see more planning within and across the economy, and what impact will this have on economic institutions?

This paper builds upon recent contributions by Phelan and Wenzel (2023) and Lambert and Fegley (2023) who argue that advances in computing technology including AI, big data, quantum computing, and the like will not resolve the economic problem as set out by Ludwig von Mises (1920 [1990], 1922 [1981]) and Friedrich Hayek (1935, 1940, 1945). Both sets of authors assert – correctly – that these advances in computing technology do not resolve the discovery problem at the heart of economic calculation.

This paper's contribution is to spell out the relationship between AI and the knowledge problem at both the more macro level (i.e. is central planning now viable?) and at the business level (what impact will AI have on planning at an organisational level?). It is often suggested that there is a disconnect between planning at the macro level and the business level, but as Klein (1996) suggests, Murray Rothbard (1962 [1970]) had proposed an answer to that question, as had Oliver Williamson (1975, 1985). Central planning within an organisation is only 'viable' in the presence of external markets that generate price signals. AI does not change that insight.

The argument made here is that the economic-knowledge problem is pervasive and, in part, manifests itself in bounded rationality. AI can go some way to resolving information problems but cannot resolve contextual knowledge problems (see Kiesling, 2015, 2023; Thomsen, 1992). Advances in computer technology do not make central planning any more viable (see Davidson, 2023 for a discussion of blockchain technology in the context of economic calculation). Bounded rationality as understood by Simon, Gigerenzer, and Smith – and operationalised by Oliver Williamson – revolves around the use of heuristics under conditions of uncertainty where local conditions and local knowledge are important considerations. Ludwig von Mises and Friedrich Hayek have explained that it is information costs that are fatal to central planning. Bounded rationality, as understood by Simon, is only partially resolved by AI. The knowledge problem as understood by Mises and Hayek is not resolved by AI. While AI is a useful and valuable tool, it does not make central planning a viable economic system at either the macro or the business level.

The paper is structured as follows: In section 'Bounded rationality: concepts', the concept of bounded rationality and its implications for economic decision-making is discussed. There are three different interpretations of the term 'bounded rationality': Herbert Simon's original meaning, a 'neoclassical' interpretation, and a behavioural interpretation. AI can resolve the neoclassical interpretation of bounded rationality. In theory it should be able to overcome the behavioural interpretation of bounded rationality, but in practice it does not. The core argument of this paper is that AI cannot and does not resolve bounded rationality problems as it was originally defined by Simon.

The section 'The knowledge problem: Mises and Hayek' examines the knowledge problem as described by Mises and Hayek and considers whether AI can overcome planning problems and make central planning a viable option. This section also contains a discussion of delegated authority and how choices can be made and improved upon through the adoption of various decision-making rules and the like. AI is a tool but does not resolve the knowledge problem as articulated by Mises and Hayek. Section 'AI and organisation: Williamson and Rothbard' re-examines the planning problem but from the perspective of Murray Rothbard and Oliver Williamson. Corporate planning is viable in the presence of external market prices and internal communication strategies that reveal local knowledge. AI in this context is an additional tool that can be used to plan, but it still does not resolve bounded rationality problems. Here the argument is that, as a result of AI, we might see more planning but we're not necessarily going to observe better planning. The question, 'What sort of technology is AI?' is examined in section 'What sort of technology is AI?'. A conclusion follows.

### Bounded rationality: concepts

Herbert Simon, one of the original founders of AI as a discipline, the 1975 Turing Award winner and 1978 economics Nobel laureate, fundamentally challenged a pivotal assumption in neo-classical economic theory with his notion of 'bounded rationality.'

Neo-classical economic theory posits that human agents are 'perfectly' rational, capable of accessing all necessary information, evaluating every potential option, and consistently making decisions that

maximise their known utility functions. Simon (1955), however, introduced a nuanced alternative view; arguing that human decision-making is inherently limited by cognitive constraints and the finite amount of information, and time available, to make the decision. Simon argued that humans are not able to find the ‘optimal solution’ to every problem, because they have limited information, limited cognitive abilities, and limited time. Therefore, humans use heuristics, or rules of thumb, to simplify the problem and find a satisfactory solution, rather than an optimal solution.

Conlisk (1996: 686) has summarised the notion of bounded rationality as follows: ‘Human cognition is a scarce resource, implying that deliberation about economic decisions is a costly activity’. He makes compelling arguments as to why bounded rationality should be included in economic models and methodology more generally. Unfortunately, however, there are differences of opinion amongst economists as to what bounded rationality is, and how to best incorporate it into economic thinking.

Gerd Gigerenzer (2017) – a psychologist who has studied the use of bounded rationality and heuristics in decision making – has argued that there are three different understandings of the term ‘bounded rationality’ in the economics literature. The first understanding is that as first proposed by Herbert Simon. He studied how people make rational decisions in real-world situations of *uncertainty*, as opposed to well-defined situations of *risk* (see Knight, 1921 for the seminal discussion of the differences between uncertainty and risk). Simon encouraged economists to study actual decisions under uncertainty rather than constructing *as if* models of expected utility maximisation.

The second understanding can be described as being ‘optimisation under constraint’. According to Gigerenzer this is how (some) economists adopted the term ‘bounded rationality’ but then subverted it to mean precisely the opposite to what Simon had intended. This approach assumes that agents have perfect information and unlimited time, but face some constraints such as budget, technology, or cognitive limitations. The decision makers then optimise their utility or profit subject to those constraints. AI is very likely to resolve bounded rationality as defined here. Conlisk (1996: 683–686), however, considers and debunks eight separate arguments that are routinely proposed in support of treating bounded rationality as simply being just another constraint to be included in an optimisation model.

The third understanding relates to how proponents of behavioural economics have appropriated the term bounded rationality. Gigerenzer argues that behavioural economists changed the meaning of bounded rationality to imply violations of rational choice theory, which they then interpreted as irrationality. This approach focuses on identifying and correcting the systematic errors or biases that people make when they use simple heuristics or rules of thumb to make decisions. Here Gigerenzer is being critical of the 2002 economics Nobel laureate Daniel Kahneman, and Amos Tversky. Gigerenzer (2015) is critical of the notion of libertarian paternalism. Similarly, Berg and Davidson (2017) argue that libertarian paternalism suffers from a knowledge problem that makes it impossible for paternalists to second-guess consumer preferences.

In theory, AI could work well to overcome human biases by providing objective analysis and decision-making based on data, rather than on human intuition, or preconceptions. In practice, however, AI has not always been particularly successful in this area (Manyika *et al.*, 2019). For example, in 2018, it was reported that an AI tool used by Amazon for recruitment had taught itself to favour male candidates over female candidates for technical jobs because it was trained on resumes submitted to the company over a 10-year period, which were predominantly from men. In 2016, it was reported that an AI tool was more likely to identify African Americans as being higher risk of recidivism than they were.

Gigerenzer’s own research follows Simon’s original definition of bounded rationality and is based on the idea that humans use simple heuristics, or rules of thumb, that are adapted to the structure of their environment and enable them to make ‘fast and frugal’ decisions under uncertainty, and time pressure. Heuristics, according to Gigerenzer, are not inferior or error-prone shortcuts, but rather smart and effective strategies that make best use of the limited information and computational capacity available to the decision maker. Gigerenzer argues that these heuristics are not just about computational efficiency, but also about robustness and simplicity in uncertain environments (Gigerenzer and Gaissmaier, 2011).

Importantly, for our purposes Gigerenzer sets out three theoretical guidelines for thinking about bounded rationality: Take uncertainty seriously. Take heuristics seriously. Take ‘ecological rationality’ seriously. This last point requires some additional explanation.

Gigerenzer *et al.* (1999: 5) defined ecological rationality as rationality that is consistent with reality. Later they provide a better definition (1999: 13), ‘A heuristic is ecologically rational to the degree that it is adapted to the structure of an environment ...’. Vernon Smith (2008: 36), who shared the 2002 economics Nobel Prize with Daniel Kahneman, paraphrases the Gigerenzer *et al.* definition as follows:

The behavior of an individual, a market, an institution, or other social system involving collectives of individuals is ecologically rational to the degree that it is adapted to the structure of its environment.

Smith (2008: 2) had already explained:

Ecological rationality refers to emergent order in the form of practices, norms, and evolving institutional rules governing action by individuals that are part of our cultural and biological heritage and are created by human interactions, but not by conscious human design.

Smith’s argument here is Hayekian. Decisions are embedded within evolved institutions, consequently context, time, and place all matter when evaluating ‘rationality’. Mousavi (2017: 91) has summarised the overlap in Gigerenzer’s and Smith’s understanding of ecological rationality and can be paraphrased as follows: Both use the same definition of ecological rationality; heuristics can be substituted by markets, firms, and rule systems; rules emerge from social behaviour and the choice of heuristic is not always deliberate; the evaluation of heuristics must be grounded in real-world experiences rather than theoretical sophistication. It is this last point that is important for our purposes – AI constitutes theoretical sophistication, but it is an open question as to how AI gains ‘real-world experience’.

Lynne Kiesling (2015: 53) describes the same phenomenon as the ‘contextual knowledge problem’:

Here I define contextual knowledge as including tacit knowledge (knowledge relevant in specific contexts that we do not know consciously that we know or how we acquired the knowledge), inarticulate knowledge (unexpressed or unspoken knowledge underlying an action or decision), and emergent knowledge that only exists in the specific context of a purposeful action or interaction.

Kiesling argues that this form of knowledge forms part of individual human perceptions and does not exist independently of a market context (see also Thomsen, 1992 for an extended discussion on this point). Importantly, Kiesling (2015: 55) argues that this knowledge cannot be replicated by an ex-ante, non-market mechanism.

AI is an ex-ante, non-market mechanism.

The point here is that AI may have access to information (many data points) but cannot have access to knowledge. As Boettke (2002: 266) explains, ‘information is the stock of the existing known, while knowledge is the flow of new and ever expanding areas of the known’. Furthermore, that new knowledge comes into existence ‘only because of the context in which actors find themselves acting’ (Boettke, 2002: 267). Information can be discovered through search functions and AI is very good at collating information. Knowledge, however, is discovered through purposeful human action. Knowledge requires interpretation and judgement, while information is basic facts (Klein, 2012: 320).

### The knowledge problem: Mises and Hayek

If solving the ‘economic problem’ were simply a matter of pure logic – as the neoclassical school suggests – then the advent of AI is hugely advantageous in that regard. An economic agent that is neither

opportunistic, nor suffers from bounded rationality (as defined by neoclassical economists), would ensure that ‘technosocialism’ could work (See Boettke and Candela, 2022 for a discussion of technosocialism). As Ludwig von Mises (1920 [1990], 1922 [1981], 1944 [2007], 1949 [1996]) and Friedrich Hayek (1935, 1940, 1945) argued, however, solving the economic problem is *not* one of pure logic (see Boettke, 2019 for a discussion of the similarity of Mises’ and Hayek’s views and analysis). As Gigerenzer has argued bounded rationality is not just another constraint that can be resolved via technology, nor is it irrationality.

In his 1922 analysis of socialism and central planning, Mises (1922 [1981]: 101) includes this passage:

But no single man, be he the greatest genius ever born, has an intellect capable of deciding the relative importance of each one of an infinite number of goods of higher orders. No individual could so discriminate between the infinite number of alternative methods of production that he could make direct judgments of their relative value without auxiliary calculations. In societies based on the division of labour, the distribution of property rights effects a kind of mental division of labour, without which neither economy nor systematic production would be possible.

Following on from the discussion of bounded rationality, it is clear that Mises has identified a similar cognitive problem. He also appears to argue that the division of labour – both physically and mentally – is a response to that condition. In his 1944 discussion of *Bureaucracy*, Mises restates that point slightly differently:

The problem to be solved in the conduct of economic affairs is this: There are countless kinds of material factors of production, and within each class they differ from one another both with regard to their physical properties and to the places at which they are available. There are millions and millions of workers and they differ widely with regard to their ability to work. Technology provides us with information about numberless possibilities in regard to what could be achieved by using this supply of natural resources, capital goods, and manpower for the production of consumers’ goods.

That last sentence could very easily be interpreted as suggesting that technology is the solution to the economic problem. Mises could not have imagined the advances in AI that have been achieved since he wrote those words. Yet – it is clear from his very next sentences that ‘technology’ is not the solution to the economic problem.

Which of these potential procedures and plans are the most advantageous? Which should be carried out because they are apt to contribute most to the satisfaction of the most urgent needs? Which should be postponed or discarded because their execution would divert factors of production from other projects the execution of which would contribute more to the satisfaction of urgent needs?

Solving the economic problem involves having humans make choices. Of course, there are many tools that can be used to inform decision making or make better decisions. For example, linear programming seems an obvious tool to answer some of the questions that Mises is asking (Dorfman *et al.* (1958)). Then there are various decision-making rules that can be employed such as maximising expected value or the minimax regret rule, and so on. Similarly, there are processes (usually labelled ‘strategic planning’) that humans can employ to improve decision making. Over time various institutions have evolved to improve decision-making – here, for example, we see the separation of powers within political governance or the separation of decision management and decision control in corporate governance (Fama and Jensen, 1983).

It is true that humans can program machines to perform tasks *as if* they were making decisions. Consider the trivial case of a vending machine: in example for a payment the vending machine provides a good. The machine itself, however, has made no choice or decision – it has been programmed by a human to perform in a particular way. All choices and decisions within an economy are ultimately made by humans. This includes delegated decisions. One of the benefits of AI – and industrialisation generally – is that many tasks previously performed by humans are now being delegated to machines.

Delegation is common within organisations (see Foss and Klein, 2012 for a discussion on delegated entrepreneurship). Delegating to a human agent, however, gives rise to the agency problem (Jensen and Meckling, 1976). Delegating to a non-human agent, such as AI, however, should not give rise to an agency problem (see Allen *et al.*, 2023). That insight, however, is trivial. Machines are simply not opportunistic (as defined by Williamson, 1975, 1985 – see below). The challenge in delegating to non-human agents is in communicating tacit knowledge (itself difficult with human agents), maladaptation (when real world conditions deviate substantially from expected conditions – see Aoki, 1983), and ensuring that the non-human agent performs as expected (the program is bug free). Delegating decision making to AI promises to economise on human cognition (see Berg *et al.*, 2018). As Alfred Whitehead indicated, ‘Civilization advances by extending the number of important operations which we can perform without thinking about them’ (quoted in Hayek, 1945: 528). It is not immediately clear that the various costs (and benefits) of AI delegation have been fully evaluated. At the very least, agency costs (delegation to another human) need to be traded off against the costs of delegating to an AI – those costs would include (net) tacit information communication costs, (net) maladaptation costs, and coding error costs. A full discussion of these issues is well beyond the scope of this paper and is a topic for further research (see Berg *et al.*, 2023 for preliminary research on these issues).

It has long been argued that computing technology can resolve the knowledge problem. Oskar Lange (1967), for example, had proposed a computing solution (see Lavoie, 1985 [2016] generally, and specifically Davidson, 2023 for a discussion of the Chilean Cybersyn project that attempted to implement this very ‘solution’). Here we discuss the 1973 economics laureate Wassily Leontief’s argument for ‘smart machines’ (quoted in Silk, 1983). The late Don Lavoie – an economist who had an undergraduate qualification in computer science – responded to this argument as follows (1985 [2016]: 118–119):

But all of this fundamentally misconceives the nature of computers. As any programmer will agree, the computer is not a smart machine. It ‘knows’ only what it is explicitly told by its programmer, who in fact is the one who does all the real mental work. ...

It is highly significant that Leontief mistakes computers for intelligent minds. If one’s view of knowledge is restricted to explicit numerical data and if one supposes mental processes to be no more than mechanical data processing, it is not surprising that one may feel a threat at being replaced by smart machines. What Leontief fails to realize is that the intellectual processes of the mind and the market processes of human societies are both undesigned and complex spontaneous orders, while the operation of the computer is a designed and relatively simple product of human minds.

While it is the case that LLMs do not need to be ‘explicitly told’ by their programmers, the general principle remains the same. AI – smart machines – only ‘know’ data that can be programmed into, or read, by a machine. Despite their ‘learning’ even LLMs, unlike humans, are not capable of ‘generalizing beyond their pretraining data’ (Yadlowsky *et al.*, 2023: 2).

It is here that we introduce Hayek (1945) into the analysis. He had argued (1945: 519–520):

The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in

concentrated or integrated form but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess. The economic problem of society is thus not merely a problem of how to allocate 'given' resources – if 'given' is taken to mean given to a single mind which deliberately solves the problem set by these 'data'. It is rather a problem of how to secure the best use of resources known to any of the members of society, for ends whose relative importance only these individuals know. Or, to put it briefly, it is a problem of the utilization of knowledge which is not given to anyone in its totality.

Hayek too could never have imagined the advances in AI technology that have occurred since he wrote those words. Yet if we examine the 'data' that are required to centrally plan, it quickly becomes clear that AI is not the solution to the economic problem either. Hayek (1945: 519) had argued, 'If we possess all the relevant information, *if* we can start out from a given system of preferences, and *if* we command complete knowledge of available means' the economic problem can be solved by logic. Hayek was responding to a claim by Oskar Lange (1936) that three pieces of information were necessary to solve the economic problem: Prices, preferences, and resources. Lange had argued that central planners with their knowledge of preferences and resources could solve for prices in a centrally planned economy. Hayek's argument was that Lange had assumed the answer to the knowledge problem. Hayek's argument is that preferences and resources are discovered; they are not 'given'. Kiesling (2023) explains this point with an excellent example of contextual knowledge.

If I asked you, out of the blue, what's your preference/willingness to pay for a can of LaCroix pamplemousse sparkling water (the best flavor!), how would you answer? Would you even know how to answer? Probably not, because knowledge like that is really only created in a context where we have to make a choice: I'm at the airport, and I can either fill my water bottle at the fountain at zero price or buy a can of LaCroix for \$1.99. Or, I've just walked 5 miles along the lakefront in Chicago and roll up at the Diversey golf driving range's cafe, and it's 80 degrees out and sunny, do I pay the \$1.49 for the can of LaCroix? That knowledge only emerges within the very personal and local context of my situation and my preferences in that situation.

It is true that AI (and digital technology generally) can be used for forecasting demand and inventory management and the like. Tools to perform these tasks already exist but it is likely that AI can drive costs savings and efficiencies in performing these business functions. It is easy to imagine that an AI that has access to localised weather information and past consumer behaviour might forecast that demand for cold drinks might be higher on some days and not others. Humans are able to make those judgements too.

It is here that Gerd Gigerenzer's notion of ecological rationality echoes Hayek's understanding of the economic problem. In Hayek's view, individuals, with their own limited and often disparate bits of knowledge, make decisions based on their local circumstances. This decentralised decision-making process reflects ecological rationality in that individuals use the knowledge and resources available in their environment to make decisions that are 'good enough' rather than seeking an optimal solution which may be practically impossible due to the dispersed nature of information.

Tacit knowledge, however, must be transformed (at some cost) into explicit knowledge and then communicated to the AI. Discovering resources that can be deployed for economic activity is an entrepreneurial insight. Israel Kirzner (1997: 34–35) describes entrepreneurial insight as follows:

The entrepreneur who 'sees' (discovers) a profit opportunity, is discovering the existence of a gain which had (before his discovery) not been seen by himself or by anyone else. Had it been seen previously, it would have been grasped or, at any rate, it would have been fully expected and would no longer then be a fresh discovery made *now*. When the entrepreneur discovers a profit opportunity, he is discovering the presence of something hitherto unsuspected.

Again, ecological rationality is important. Entrepreneurs make decision based on local knowledge and preferences in an environment where they believe that prices are ‘wrong’. Entrepreneurs need to understand the context in which they operate. This requires tacit knowledge, or knowledge gained from personal experience. AIs are trained on existing (given) information sets. They make probabilistic decisions i.e. they are able to make decisions involving risk, not uncertainty. In short, AIs are backward looking where entrepreneurs are forward looking.

Potts (2010, 2017) makes the argument that even for humans’ entrepreneurial insight is difficult. He identifies ten separate behavioural characteristics that exacerbate what he calls ‘the problem of choice over novelty under uncertainty’. The very first characteristic is ‘awareness of novelty is hard’. Potts’ argument is that in a world of rapid change that humans tend to overlook genuine novelty. If humans struggle to observe novelty under real-world conditions, it is unclear how AI that is entirely backward looking can observe novelty.

Phelan and Wenzel (2023: 177–178) make a similar point. They argue that to solve the economic problem, ‘we would need an artificially intelligent entrepreneur’. They go on to explain that what would be required is, ‘a machine that can ‘think outside the box’’. The problem with that argument – although Phelan and Wenzel do not use this terminology – is that AI cannot ‘think’ *outside* the box; the AI *is* the box. Similarly, Lambert and Fegley (2023: 246) point out that all the data used to train AI, ‘are entirely that of the past’. AI that is trained on a sub-set of known information that can be easily communicated in machine readable form cannot ‘know’ enough to replicate the entrepreneurial function.

### AI and organisation: Williamson and Rothbard

New institutional economics (NIE) is a branch of economics that studies how institutions shape human behaviour and economic outcomes. Institutions are the formal and informal rules, norms, and customs that govern social interactions, such as laws, contracts, property rights, markets, organisations, and culture. NIE assumes that individuals are rational and self-interested, but also bounded by cognitive limitations, incomplete information, and transaction costs. Therefore, institutions matter because they reduce uncertainty, facilitate cooperation, and coordinate expectations among economic agents. Economists who have made significant contributions in this field include Armen Alchian (2006), Yoram Barzel (1989, 2002), Ronald Coase (1937, 1960), Harold Demsetz (1988), Douglass North (1990), Elinor Ostrom (1990), and Oliver Williamson (1975, 1985) (see also Foss and Foss, 2022: 37–62 for further discussion).

Oliver Williamson, the 2009 Nobel laureate in economics (joint with Elinor Ostrom), argued that bounded rationality affects the choice of governance structures, because it exacerbates uncertainty and complexity in transactions. For example, bounded rationality leads to incomplete contracts that cannot specify all possible contingencies. Bounded rationality may also limit the ability of parties to monitor and enforce contracts.

According to Williamson (1985), bounded rationality is one of the two main behavioural assumptions of Transaction Costs Economics (a sub-field within NIE); opportunism is the other behavioural assumption. Opportunism is the tendency of some actors to not only act in their own self-interest but pursue their self-interest with guile. Williamson claimed that opportunism is a realistic assumption that reflects human nature and cannot be ignored in economic analysis. He also suggested that opportunism is exacerbated by bounded rationality, because it contributes to information asymmetry, leading to adverse selection and moral hazard.

Bounded rationality is important to Oliver Williamson’s research agenda because it explains why different governance structures emerge and how they affect economic performance. Governance – how Williamson describes ‘hierarchy’ – emerges when bounded rationality, opportunism, and asset specificity are present. In the absence of asset specificity, we see what Williamson labels ‘competition’ i.e. market transactions. In Williamson’s schema, the absence of bounded rationality, or unbounded rationality as he labels it – gives rise to ‘Planning’. Planning is an efficient contracting process in the context of unbounded rationality.

Again, this invites speculation that AI can reduce bounded rationality problems and can facilitate planning. It is here that the disconnect between central planning at a national level and planning at a corporate level becomes important. Large corporations appear to be very successful at planning – yet central planning at a national level is unsuccessful. Ronald Coase had asked the question, ‘Why is not all production carried on in one big firm?’. Some multinational corporations are ‘larger’ than some nation-states – how is it that these corporations appear to plan successfully, but those smaller nation-states cannot? What are the limits to corporate planning and does AI change those limits?

Klein (1996: 7) argues that several answers to Coase’s question have been suggested, including some by economics laureates:

Existing contractual explanations rely on problems of authority and responsibility (Arrow, 1974); incentive distortions caused by residual ownership rights (Grossman and Hart, 1986; Holmstrom and Tirole, 1989); and the costs of attempting to reproduce market governance features within the firm (Williamson, 1985, chap.6).

Klein suggests that Rothbard (1962 [1970]: 544–560) has a better solution to Coase’s question that involves extending Mises’ economic calculation problem to the firm. Rothbard makes the argument that vertically integrated firms are only able to establish profitability by reference to prices that are determined in external markets – *‘there can be no implicit estimates without an explicit market!’* (1962 [1970]: 543, emphasis original). In Rothbard’s framework, this creates an upper bound on the size of the firm – for every capital good (stage of production in Rothbard’s explanation) there must be a market where that good is traded and a market price established (1962 [1970]: 548):

Because of this law, firms cannot merge or cartelize for complete vertical integration of stages or products. Because of this law, there can never be One Big Cartel over the whole economy or mergers until One Big Firm owns all the productive assets in the economy. ... As the area of incalculability increases, the degrees of irrationality, misallocation, loss, impoverishment, etc., become greater. Under *one* owner or *one* cartel for the whole productive system, there would be no possible areas of calculation at all, and therefore complete economic chaos would prevail.

Here Rothbard has not only reconciled Mises’ and Hayek’s views with those of Coase, but he has also provided an additional insight. Planning fails when it is centralised. Successful planning is local planning that takes into account local knowledge. In his recent history of the modern corporation over the twentieth century, Richard Langlois (2023: 152) provides (anecdotal) evidence to support Rothbard’s argument.

GM [General Motors] put in place a variety of controls, including controls over capital allocation and inventories. Divisions would also be evaluated using Donaldson Brown’s momentous financial innovation, the principle of return on investment. Even daily cash flows would be managed, through an innovative system of interbank clearing that the Federal Reserve was experimenting with. In all these cases, local information was being codified so that it could be transmitted to the center. Despite this effort, however, a downturn in 1924 revealed that unlike Ford, GM did not actually have reliable data on the sales of its cars. [Alfred] Sloan, who had been elevated to the presidency of the corporation in May 1923, hurriedly created a system of reporting by dealers and hired an outside firm to supply data on new car registrations.

General Motors had attempted to centralise decision-making and codify local information to facilitate that centralisation. Yet without external market knowledge it was unable to manage its internal operations and planning efforts. Langlois (2023: 151) summarises this insight as follows: ‘Federal decentralization – the M-form – works well when it is possible to identify business units that produce

a distinct product for a distinct market.’ That is Rothbard’s argument too. External market information is necessary to facilitate internal planning.

Williamson (1985: 52) too is hostile to the notion of central planning, but not just on the grounds of bounded rationality.

But utopian societies are especially vulnerable to the pound of opportunism.

The new man of socialist economics is endowed with a high level of cognitive competence (hence the presumed efficacy of planning) and displays a lesser degree of self-interestedness (a greater predisposition to cooperation) than his capitalist counterpart.

His major argument against central planning (note that central planning is not his research interest) is that central planners are opportunistic. This differs from Mises and Hayek who were abstracting from the ‘incentive problem’. Williamson’s argument does not explicitly rule out a technology that reduces bounded rationality making central planning (more) viable. It is important to note, however, that Williamson’s arguments do not exclude strategic planning within organisations. What is particularly important for our purposes is that Williamson explains the emergence of the M-form organisation as a form of decentralisation of decision making.

In his 1975 treatment of markets and hierarchies, Williamson (1975: 4–5) carefully distinguishes his position from Hayek’s position on central planning. After summarising Hayek’s (1945) analysis, Williamson observes (emphasis original):

Although each of these observations is important to the argument of this book, I use them in a somewhat different way than does Hayek – mainly because I am interested in a more microeconomic level of detail than he. Given bounded rationality, uncertainty, and idiosyncratic knowledge, I argue that prices often do *not* qualify as sufficient statistics and that a substitution of internal organization (hierarchy) for market-mediated exchange occurs on this account.

It appears that Williamson’s (1975) rejection of the possibility of central planning is even stronger, indeed more emphatic, than that of Mises and Hayek. Mises (1920 [1990], 1922 [1981]) had argued that in the absence of money prices that economic calculation was impossible. Even in the presence of money prices, Williamson argues that central planning will fail. That insight informs his interpretation of the emergence of the M-form organisation. Williamson bases his argument on Chandler’s (1962, 1977) discussion of the emergence of the M-form at General Motors. Langlois (2023), building on Freeland (2001), suggests that the actual functioning on the M-form organisation at GM was somewhat different to how it has been described. It is not immediately clear, however, that the historical discrepancy is (sufficiently) inconsistent with Williamson’s theory.

Williamson discusses bounded rationality in far more detail in his 1975 book than he does in 1985. His 1985 treatment seems perfunctory. Using Gigerenzer’s language it seems that Williamson (1985) is taking uncertainty seriously, but it is not clear that he is taking heuristics seriously (see Foss, 2001, 2003), nor seemingly does he incorporate ecological rationality into his analysis. His 1975 treatment of bounded rationality, however, is far more complete.

Williamson (1975: 21–23) contains his most comprehensive discussion on bounded rationality. There he explains that bounded rationality includes both ‘neurophysiological limits’ and ‘language limits’. The neurophysiological limits relate to limitations on individuals’ ability to ‘receive, store, retrieve, and process information without error’ (1975: 21). Language limits are limitations on the ability of individuals to communicate with other individuals in a manner that fully describes their intent (1975: 22). It seems plausible, however, that AI can relax a bounded rationality constraint if it is defined in these terms. AI can outperform human intelligence in storing, retrieving, and processing information. Similarly, AI’s can communicate with other AI’s in unambiguous terms.

It is here that Williamson (1975: 22) discusses ecological rationality – while not actually using that terminology.

Bounds on rationality are interesting, of course, only to the extent that the limits of rationality are reached – which is to say, under conditions of uncertainty and/or complexity. In the absence of either of those conditions, the appropriate set of contingent actions can be fully specified at the outset. Thus, it is bounded rationality in relation to the conditions of the environment that occasions the economic problem.

Surprisingly, Williamson does not differentiate between risk and uncertainty in a Knightian sense. He does, however, argue that uncertainty and complexity is exacerbated by ‘information impactedness’. This is not just a situation of asymmetric information, but an information problem that cannot be costlessly resolved (if at all). In his later work (1985: 21) he also refers to ‘maladaptation costs’ as being a source of uncertainty. This is the cost of having to renegotiate contracts after changing circumstances have altered the intended outcome of the initial contract.

Williamson describes the emergence of the M-form organisation in similar terms to Rothbard’s explanation of why a single firm does not dominate the entire economy. Profit centres are established where profitability can be observed (presumably with reference to external markets) and operational decisions are made at lower levels of the hierarchy while strategic decisions are made at higher levels of the hierarchy. As Freeland (2001) and Langlois (2023) have argued, historically there was not a clear-cut differentiation between strategic and operational decision making. As Langlois (2023: 153) explains:

What Sloan implemented on the ground was instead ‘a consultative style of top management in which authority became more firmly tied to technical expertise.’ ... Top executives could not just tell subordinates what to do. They had to provide reasons and marshal facts to back up their demands, and subordinates in turn could draw on their own local expertise to contest decisions. Far from being decoupled, knowledge at the center and in the divisions would interact. In essence, Sloan was attempting to create an organization that could learn in much the same way that scholars – at their best, at least – learn through contesting one another’s conjectures.

What Langlois is describing here is a discovery process where local knowledge is transformed into general knowledge (at least within the organisation). This is not a central planning process where local information is codified and transmitted to the centre. Rather this is a discovery process where in the face of actual decisions being made, decision makers must ‘actively interpret the information [they] receive, and pass judgement on its reliability and its relevance for ... decision-making’ (Boettke, 2002: 267). This is a process where information is transformed into the knowledge necessary to make decisions.

These issues that Williamson identifies are unlikely to be resolved by AI. How AI ‘receives’ information is important to resolving bounded rationality problems. Once information is received AI can outperform a human at many margins. Communicating that information in a digitally readable format, however, remains a challenge and AI faces a ‘language’ challenge (albeit of a different character) as do humans.

The fact remains that AI can overcome many computational limitations that humans face and as such relax some of the bounded rationality constraints within Williamson’s framework. The language problem of communicating information to the AI, however, remains. AI does not fully eliminate bounded rationality. It does, however, expand the range of planning that can occur. In the Williamson framework, that observation leads to the prediction that we should observe *less* hierarchy and *more* planning because of AI. That observation, however, does not necessarily imply *better* planning.

It can be argued that the increase in planning could circumvent the bounded rationality problem. Williamson is ultimately concerned with the problem of incomplete contracts. AI could provide a ‘brute force solution’ to incomplete contracting by simply simulating all (or at least, many) possible

future scenarios and then developing a contingency plan for each situation. Kiesling (2023), for example, speculates as to this very situation:

We can analyze patterns in data after the fact, and we can use AI and machine learning to help us make out-of-sample predictions of future outcomes using those historical data ...

As she indicates this ‘solution’ does not resolve the contextual knowledge problem, but it could resolve the incomplete contracting problem. As indicated above, however, at present AI is not capable of making accurate out-of-sample predictions (Yadlowsky *et al.*, 2023). That capability may improve in future but that may fall into Minsky’s argument that ‘easy things are hard’ – humans can easily imagine the future.

In summary, AI relaxes some of the constraints that bounded rationality imposes on economic activity that then gives rise organisational forms such as corporate hierarchy. It does not, however, overcome problems of information impactedness nor maladaptation of contracts to changing circumstances. AI should result in less hierarchical organisational forms, but it is not clear if that implies fewer hierarchies or smaller hierarchies.

### What sort of technology is AI?

AI is best understood as a new general-purpose technology (GPT) that has ubiquitous application and transformative potential across different sectors of the economy. Like the steam engine, electricity, computers, or the internet, which have greatly transformed both the economy and society at large, AI is not bound to a single specific application but is foundational, opening up wide arrays of uses. By automating tasks that required human cognition, from basic data processing to complex problem solving, it effectively augments human capital. This augmentation of human capital has the potential to dramatically increase productivity, mirroring the impact of previous GPTs like electricity, which extended the working day, or the Internet, which revolutionised communications and information access. Moreover, AI facilitates coordination across different areas of human activity by enabling efficient data processing, prediction, and decision-making, thereby reducing the cognitive load on individuals, and enhancing productivity in a broad array of tasks.

While AI can optimise existing institutional mechanisms that economise on bounded rationality, providing more efficient tools for coordination and decision-making, it is less likely to induce the emergence of entirely new institutional structures. AI can simplify complexity but does not resolve contextual knowledge problems. Human insight and decision-making remain necessary to resolve the economic problem at both the macro and micro levels of the economy.

### Conclusion

Artificial intelligence has the potential to profoundly disrupt the economy. With its ability to analyse vast amounts of data and make intricate predictions, it can reshape industries, making them more efficient and potentially introducing new modes of production and distribution (Agrawal *et al.*, 2018). Furthermore, through automation and optimisation, AI can reduce costs and increase productivity in various sectors, including financial services, manufacturing, logistics, and various other service industries. AI promises to significantly enhance decision-making processes by overcoming human cognitive limitations. With the potential to sift through, and analyse, vast datasets within a fraction of the time that a human would require, AI can derive valuable insights and present decision-makers with an abundance of nuanced information and decision options. Moreover, AI offers to mitigate human decision-making biases, which often cloud judgment and lead to sub-optimal decisions.

In theory, AI has the potential to overcome bounded rationality. In practice, however, it does not.

Despite its advanced computational abilities, however, AI does not make central planning a viable solution to the economic problem. This is primarily due to AI’s inherent inability to address the

knowledge problem, a concept described in detail by Ludwig von Mises and Friedrich Hayek. In essence, both humans and AI lack the ability to collate and interpret the vast, decentralised, and often subjective information distributed among individuals in an economy, underlining the continued relevance and necessity of market mechanisms. It is these very market mechanisms that makes corporate planning viable.

In conclusion, while AI is a potent tool with significant potential to drive efficiency, enhance productivity, and spur innovation within the economy, it does not provide a comprehensive solution to the complex problem of economic organisation at either the macro or the micro level. Despite its capacity to challenge traditional economic paradigms and ameliorate some of the effects of bounded rationality, AI has inherent limitations.

Despite the fact that AI can beat humans at chess or Go and solve complex math problems and the like, there remains many things that it cannot do. AI researchers are attempting to replicate tasks that humans can perform and have achieved astonishing success. But as Marvin Minsky argued, 'Easy things are hard'. Humans can easily perform a wide range of tasks that AI cannot, particularly those involving creativity and nuanced understanding of context. As Potts (2010, 2017) suggests, even humans find entrepreneurial insight difficult – at present that is impossible for AI. At present AI cannot solve 'contextual knowledge problems' – they are not ecologically rational.

Despite the astonishing advances made in AI over the past decade there is much that we do not know. As AI researchers expand the scope of what can be done, we as economists can gain a greater understanding of choice and decision-making. An immediate question is when can decision-making be delegated to an AI? Berg *et al.* (2023) argue that, at present, AI is too unreliable to be trusted for any decisions that require discretion and propose that smart contracts be deployed to constrain AI. Then there is a question as to whether AI can make decisions without any human intervention. This, in turn, raises philosophical questions – what does it mean to make a decision? How do humans make decisions? Do humans entirely have free will when making choices? Reasonable people can reasonably disagree on these issues. What is clear, however, is that neither AI researchers nor economists have satisfactory answers to these questions.

### *Declaration of generative AI and AI-assisted technologies in the writing process*

During the preparation of this work the author used Bing Chat, ChatGPT, and Google Bard (see Berg, 2023) in order to reword sentences, correct grammar, and to simplify the explanation of artificial intelligence. After using these tools/services, I have reviewed and edited the content as needed and take full responsibility for the content of the publication.

**Acknowledgements.** Blockchain Innovation Hub, RMIT University, Melbourne, Australia. I would like to thank three anonymous referees for very helpful comments that have greatly improved this paper.

### References

- Agrawal A., Gans J. and Goldfarb A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston, Massachusetts: Harvard Business Review Press.
- Alchian A. (2006). *The Collected Works of Armen Alchian*. Indianapolis: Liberty Fund.
- Allen D., Berg C. and Davidson S. (2020). *The New Technologies of Freedom*. Great Barrington: The American Institute for Economic Research.
- Allen D., Berg C., Ilyushina N. and Potts J. (2023). Large Language Models Reduce Agency Costs. *SSRN Working Paper*. Available at [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4437679](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4437679)
- Aoki M. (1983). Managerialism revisited in the light of bargaining-game theory. *International Journal of Industrial Organization* 1(1), 1–21.
- Arrow K. (1974). *The Limits of Organization*. New York: WW Norton & Company.
- Barzel Y. (1989). *Economic Analysis of Property Rights*. Cambridge: Cambridge University Press.
- Barzel Y. (2002). *A Theory of the State: Economic Rights, Legal Rights, and the Scope of the State*. Cambridge: Cambridge University Press.

- Berg C. (2023). The Case for Generative AI in Scholarly Practice. SSRN Working Paper. Available at <https://ssrn.com/abstract=4407587>
- Berg C. and Davidson S. (2017). Nudging, calculation, and utopia. *Journal of Behavioral Economics for Policy* 1(S), 49–52.
- Berg C., Davidson S. and Potts J. (2018). Beyond Money: Cryptocurrencies, Machine-Mediated Transactions and High Frequency Bartering. Available at SSRN: [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3158047](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3158047)
- Berg C., Davidson S. and Potts J. (2023). Institutions to constrain chaotic robots: Why generative AI needs blockchain. Available at SSRN: <https://ssrn.com/abstract=4650157>
- Boettke P. (2002). Information and knowledge: Austrian economics in search of its uniqueness. *The Review of Austrian Economics* 15, 263–274.
- Boettke P. (2019). How Misesian was the Hayekian research program? *Procesos De Mercado: Revista Europea De Economía Política* 16(1), 251–257.
- Boettke P. and Candela R. (2022). On the feasibility of technosocialism. *Journal of Economic Behavior and Organization* 205, 44–54.
- Chandler A. (1962). *Strategy and Structure: Chapters in the History of the Industrial Empire*. Cambridge: MIT Press.
- Chandler A. (1977). *The Visible Hand: The Managerial Revolution in American Business*. Cambridge: Belknap Press.
- Coase R. (1937). The nature of the firm. *Economica* 4(16), 386–405.
- Coase R. (1960). The problem of social cost. *Journal of Law and Economics* 3, 1–44.
- Conlisk J. (1996). Why bounded rationality? *Journal of Economic Literature* 34(2), 669–700.
- Davidson S. (2023). Blockchain and the information – calculation problem. *Journal of Economic Behavior and Organization* 213, 142–150.
- Demsetz H. (1988). *The Organization of Economic Activity* (2 volumes). London: Blackwell.
- Dorfman R., Samuelson P. and Solow R. (1958). *Linear Programming and Economic Analysis*. New York: McGraw-Hill.
- Fama E. and Jensen M. (1983). Separation of ownership and control. *The Journal of Law and Economics* 26(2), 301–325.
- Foss N. (2001). Bounded rationality in the economics of organization: Present use and (some) future possibilities. *Journal of Management and Governance* 5, 401–425.
- Foss N. (2003). The rhetorical dimensions of bounded rationality: Herbert A. Simon and organizational economics. In Rizzello S. (ed.), *Cognitive Paradigms in Economics*. Milton Park: Routledge, pp. 158–176.
- Foss K. and Foss N. (2022). *Economic Microfoundations of Strategic Management: The Property Rights Perspective*. London: Palgrave MacMillan.
- Foss N. and Klein P. (2012). *Organizing Entrepreneurial Judgment: A New Approach to the Firm*. Cambridge: Cambridge University Press.
- Freeland R. (2001). *The Struggle for Control of the Modern Corporation: Organizational Change at General Motors, 1924–1970*. Cambridge: Cambridge University Press.
- Gigerenzer G. (2015). On the supposed evidence for libertarian paternalism. *Review of Philosophy and Psychology* 6, 361–383.
- Gigerenzer G. (2017). What is bounded rationality? In Viale R. (ed.), *Routledge Handbook of Bounded Rationality*. Milton Park: Routledge, pp. 55–69.
- Gigerenzer G., Todd P. and the ABC Research Group (1999). *Simple Heuristics That Make Us Smart*. Oxford: Oxford University Press.
- Gigerenzer G. and Gaissmaier W. (2011). Heuristic decision making. *Annual Review of Psychology* 62, 451–482.
- Grossman S. and Hart O. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy* 94(4), 691–719.
- Hayek F. (1935). Present state of the debate. In Hayek F. (ed.), *Collectivist Economic Planning: Critical Studies on the Possibilities of Socialism*. London: Routledge, pp. 201–243.
- Hayek F. (1940). Socialist calculation: The competitive ‘solution’. *Economica* 7(26), 125–149.
- Hayek F. (1945). The use of knowledge in society. *The American Economic Review* 35(4), 519–530.
- Holmstrom B.R. and Tirole J. (1989). The theory of the firm. In Schmalensee R. and Willig R. (eds.), *Handbook of Industrial Organization*. Amsterdam: Elsevier, pp. 61–133.
- Jensen M. and Meckling W. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics* 3(4), 305–360.
- Kiesling L. (2015). The knowledge problem. In Boettke P. and Coyne C. (eds), *The Oxford Handbook of Austrian Economics*. Oxford University Press, pp. 45–64.
- Kiesling L. (2023). *Markets Are Knowledge Ecosystems*. Knowledge Problem. Available at <https://knowledgeproblem.substack.com/p/markets-are-knowledge-ecosystems>
- Kirzner I. (1997). *How Markets work: Disequilibrium, Entrepreneurship and Discovery*. London: Institute of Economic Affairs.
- Klein P. (1996). Economic calculation and the limits of organization. *The Review of Austrian Economics* 9(2), 3–28.
- Klein D. (2012). *Knowledge and Coordination: A Liberal Interpretation*. Oxford: Oxford University Press.
- Knight F. (1921). *Risk, Uncertainty and Profit*. New York: Houghton Mifflin.
- Lambert K. and Fegley T. (2023). Economic calculation in light of advances in big data and artificial intelligence. *Journal of Economic Behavior and Organization* 206, 243–250.

- Lange O. (1936). *On the Economic Theory of Socialism: Part One. The Review of Economic Studies* 4(1), 53–71.
- Lange O. (1967). The computer and the market. In Feinstein C. (ed.), *Socialism, Capitalism and Economic Growth: Essays Presented to Maurice Dobb*. Cambridge: Cambridge University Press, pp. 158–161.
- Langlois R. (2023). *The Corporation and the Twentieth Century: The History of American Business Enterprise*. Princeton: Princeton University Press.
- Lavoie D. (1985 [2016]). *National Economic Planning: What is Left?* Arlington: Mercatus Center.
- Manyika J., Silberg J. and Presten B. (2019). What Do We Do About the Biases in AI. *Harvard Business Review*, October 25.
- Mises L. (1920 [1990]). *Economic Calculation in the Socialist Commonwealth*. Auburn: Mises Institute.
- Mises L. (1922 [1981]). *Socialism: An Economic and Sociological Analysis*. Indianapolis: Liberty Fund.
- Mises L. (1944 [2007]). *Bureaucracy*. Indianapolis: Liberty Fund.
- Mises L. (1949 [1996]). *Human Action: A Treatise on Economics*. San Francisco: Fox & Wilkes.
- Mitchell M. (2019). *Artificial Intelligence: A Guide for Thinking Humans*. New York: Farrar, Straus and Giroux.
- Mousavi S. (2017). Gerd Gigerenzer and Vernon Smith: Ecological rationality of heuristics in psychology and economics. In Frantz R., Chen S.-H., Dopfer K., Heukelom F. and Mousavi S. (eds), *Routledge Handbook of Behavioral Economics*. London: Taylor & Francis, pp. 88–100.
- North D. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.
- Ostrom E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.
- Phelan S. and Wenzel N. (2023). Big data, quantum computing, and the economic calculation debate: Will roasted cyberpi-geons fly into the mouths of comrades? *Journal of Economic Behavior and Organization* 206, 172–181.
- Potts J. (2010). Can behavioural biases in choice under novelty explain innovation failures? *Prometheus* 28(2), 133–148.
- Potts J. (2017). Behavioral innovation economics. In Frantz R., Chen S.-H., Dopfer K., Heukelom F. and Mousavi S. (eds), *Routledge Handbook of Behavioral Economics*. London: Taylor & Francis, pp. 392–404.
- Rothbard M. (1962 [1970]). *Man, Economy and State: A Treatise on Economic Principles*. Los Angeles: Nash Publishing.
- Russell S. and Norvig P. (2016). *Artificial Intelligence: A Modern Approach*. Indianapolis: Pearson.
- Silk L. (1983). Structural joblessness. *The New York Times*, April 6. Available at <https://www.nytimes.com/1983/04/06/business/economic-scene-structural-joblessness.html> (accessed 14 June 2023).
- Simon H. (1955). A behavioural model of rational choice. *The Quarterly Journal of Economics* 69(1), 99–118.
- Smith V. (2008). *Rationality in Economics: Constructivist and Ecological Forms*. Cambridge: Cambridge University Press.
- Thomsen E. (1992). *Prices and Knowledge: A Market-Process Perspective*. London: Routledge.
- Williamson O. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: The Free Press.
- Williamson O. (1985). *The Economic Institutions of Capitalism*. New York: The Free Press.
- Yadlowsky S., Doshi L. and Tripuraneni N. (2023). Pretraining data mixtures enable narrow model selection capabilities in transformer models. *arXiv Preprint*. Available at <https://arxiv.org/abs/2311.00871>