U.PORTO
FEP FACULDADE DE ECONOMIA
UNIVERSIDADE DO PORTO

Can AI replace the Market? A Hayekian Point of View

Pedro Miguel Santos Correia dos Reis

Dissertation
Master in Economics

Supervised by
Mário Graça Moura

September 2026

Acknowledgements

I would firstly like to thank my girlfriend Rita, without whose love, support and encouragement I would not have been able to take on the challenge of studying while working.

I would also like to thank my supervisor, Professor Mário Graça Moura. This dissertation would not have been possible without such an active and insightful supervisor. I thoroughly enjoyed our conversations over the past year.

Finally, a thank you to my family, who were by my side throughout this process.

This dissertation is dedicated to my grandmother Carolina.

Abstract:

Artificial intelligence (AI) is the defining technology of the decade. As its integration in society deepens, I investigate what is the role of markets in a world with AI. Could artificial intelligence replace the market? The debate on the possibility or impossibility of an economy without markets, at least for higher-order goods, has surfaced since the 1920s, mainly in the form of arguments on the viability of socialism. Mises and Hayek formulated thorough arguments on why an efficient economy is impossible without the market process and the prices it provides. After a century of failed experiments in socialism, artificial intelligence presents new possibilities of centralised price-setting mechanisms through increased computing power and machine learning capabilities. Despite the advances in artificial intelligence, the Austrian arguments defending markets are still valid, or so I shall argue. AI-based models are unable to discover knowledge, which is dispersed, contextual, and often tacit, as Hayek carefully explained over 80 years ago.

JEL codes: B13, B25, B40, B53, P10, P20

Keywords: Artificial Intelligence, Economic Markets, Austrian Economics, Socialist Calculation Debate, History of Economic Thought, Hayek, Mises

Resumo:

A inteligência artificial (IA) é a tecnologia marcante da década. À medida que sua integração e uso na sociedade aumentam, eu investigo qual será o papel do mercado no mundo com IA. Será que a IA poderá substituir o mercado? O debate sobre a possibilidade ou impossibilidade de uma economia sem mercados, pelo menos para os bens de ordem superior, existe desde os anos 20 do séc. XX, na forma de debates sobre o socialismo. Mises e Hayek formularam argumentos detalhados sobre o porquê de uma economia eficiente ser impossível sem o processo de mercado e os preços que este fornece. Após um século de experiências falhadas de socialismo, a inteligência artificial apresenta novas possibilidades de centralização dos mecanismos de formação de preços, através de cada vez maior poder computacional e técnicas de *machine learning*. Apesar dos avanços tecnológicos da inteligência artificial, os argumentos austríacos em defesa do mercado continuam válidos, ou assim irei argumentar. Modelos baseados em IA não são capazes de descobrir conhecimento, que é disperso, contextual e frequentemente tácito, como Hayek explicou há mais de 80 anos.

Palavras-chave: Inteligência Artificial, Mercados Económicos, Escola Austríaca, História do Pensamento Económico, Hayek, Mises

Table of Contents

1. Introduction 1

2. Artificial Intelligence and the Market 3

2.1. Concepts and Limits of AI 3

2.2. AI and the Market in 2026 5

3. From Mises to the Socialist Calculation Debate and beyond 6

3.1. Main Austrian Arguments and their Opponents up to the late 1940s 6

3.2. On recent, pre-AI Proposals of Socialism 9

3.3. New Proposals: How could AI replace the Market? 10

3.4. Austrian Proposals on AI 14

4. On the Feasibility of AI Replacement of Markets 19

4.1. Speculations on the End of Markets 19

4.2. Price-Setting under AI 21

4.3. The Nature of Knowledge 24

5. Conclusion 27

6. Bibliography 28

1. Introduction

The artificial intelligence (AI) revolution is under way worldwide, transforming the world in every domain, from computation to medicine. Investment in AI is booming: $252bn in 2024 (Kariuki, 2025), $582bn in 2025, showing no signs of slowing down (Stanford HAI, 2026). Nine of the ten most valuable firms in the world deal directly or indirectly in AI development (TradingView, 2025). The transformative impact of AI on the world economy, as well as the potential impact of future AI adoption, warrant an economic analysis of its effects on resource allocation and market competition. Will a centralised AI model be able to replicate the market functions and replace competition?

The awesome computing power of AI places the question of centralisation in a new light. Even when compared to advanced programs of the modern age, artificial intelligence-based models and agents provide far superior processing power and ease of use. Tasks and jobs considered too complex for traditional programs are increasingly automated, to the point of being achievable with a simple prompt. The question, then, is whether such a model, powerful as it may be, can manage a complex, modern economy in a more efficient way than the market hitherto has.

Economists and other social scientists have discussed the merits of markets up against centralisation, though, of course, in different technological contexts – most interestingly in the debates on the possibility or impossibility of socialism triggered by Ludwig von Mises' 1920 article. Despite the failure of self-styled 'socialist' economies in the 20th century, the topic remained alive as new technologies raised hopes of surpassing socialist shortcomings.

Mises provided a theorem on the impossibility of socialism: certain market functions, he argued, could only be performed by the market, which ruled out the rationality of an economy without markets for higher-order goods. Oskar Lange attempted a counterargument, but its flaws were made apparent by F. A. Hayek in the 'Socialist Calculation Debate', even though Hayek was understood only decades later. Hayek (1945, 1968) emphasises that market competition is a procedure to discover knowledge. Knowledge is the provisional product of a dynamic process, where ongoing interactions validate or falsify subjective 'guesses'. Prices and subjective 'guesses' trigger individual, decentralised action, which jointly leads to new knowledge (or the revision of knowledge) and new prices.

Socialism failed in practice but survives as an ideal. The advent of new information technology, namely the internet, big data and now AI, has reopened discussions, with the socialist cohort arguing that the technologies that capitalism has developed, unimaginable during the Socialist Calculation Debate, would displace the Austrian argument that a socialist economy cannot effectively convey pricing information. Of course, if that argument were right, it remains to be seen whether, under these circumstances, a government by huge corporations would not be a more likely outcome than the implementation of the socialist ideal.

My focus, however, is not on a debate on socialism. The Socialist Calculation Debate serves as a framework to debate the merits of the market through an Austrian lens and analyse whether AI is powerful and advanced enough to replace the market's ability of discovering and disseminating new knowledge, a feat no technological advancement has achieved so far.

To this end, I start by summarising the state of modern AI and the main issues of the Socialist Calculation Debate, before analysing recent proposals from both economists and tech researchers on whether and how AI could impact or even replace the market. I then scrutinise how such a replacement of the market could occur and finish with a brief conclusion.

2. Artificial Intelligence and the Market

A key element to the discussion is the current state of the technology at hand, artificial intelligence. A debate on the long-term prospects of AI is not reasonable, because of massive unpredictability as to its evolution. A comprehensive review of current technology is likewise beyond the scope of this paper. What follows is a short overview of current technology, as well as some minimal speculation on its future in as much as it is directly relevant to the argument of this paper.

2.1. Concepts and Limits of AI

Many concepts are important to understanding the argument I will make, starting with the meaning of "artificial intelligence" and the important distinction it has to large language models (LLMs). Encyclopaedia Britannica defines artificial intelligence as "the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings" (Copeland, 2026). Therefore, the term AI refers to an ability by digital devices that already exist, as opposed to a new technology in and of itself. Despite the recent developments surrounding AI and societal interest in it, the term has been used in computing since the 1950s, with studies and conferences on the intelligence of machines (McCarthy *et al.*, 1955).

The AI boom of the 2020s was sparked by a more recent form of artificial intelligence: generative AI, or gen AI. This describes AI capable of generating a reply to a user's prompt, creating something new: this can be in the form of text, media or software code, among others (Stryker & Kavlakoglu, 2024). Large language models are a form of gen AI: these are models based on statistical probability, programmed through machine learning to learn patterns and apply them in responding to user-made prompts. These models are trained on a set of information, using algorithms to predict what the desired outcome is (Stryker, 2026). The ability to digest large amounts of information drives these models' capacity to communicate with users in a deeper and less structured form than traditional programs. While a traditional program responds to user inputs in a fixed and pre-programmed way, a large language model's machine learning mechanism allows it to digest an unstructured input and guess through probabilistic algorithms what the best reply to a request is (RIGB, 2023). Machine learning allows the models to learn by interacting with the users, improving their ability to predict the desired response. The constant learning process means that LLMs

require tremendous computing power in order to function, but also that they have much more processing power than previous technologies (Bergmann, 2025). There are three types of learning for an LLM: reinforcement, supervised and unsupervised learning. Through these methods, the model improves its ability to make correct predictions (Murel & Kavlakoglu, 2026).

Agentic AI is a similar form of AI. An agent is an AI system specialised in specific tasks (Stryker, 2025). It is quite popular, especially in corporate settings, as it is more independent than generative AI. Its focus on a single task or goal allows it to specialise in solving it with less human supervision and more interaction with external tools and software. This is possible due to its reinforcement learning mechanism: the agent is programmed with a goal and positively rewarded when it achieves the goal. The trade-off is the increased risk of poor behaviours if the goal unintentionally promotes them due to programming flaws (Stryker, 2025).

These models are not only powerful, but also publicly available and easy to use. Leading LLMs, such as OpenAI's ChatGPT or Anthropic's Claude, are available for public use and present the user with a simple text box, where he or she can type in their prompt, starting a conversation with the model. The prompt can range from simple questions to complex requests, such as coding or video processing.

Despite their technological prowess, these models still fall under the scope of narrow AI. This form of AI is, simply put, one that is able to perform a task, often with tremendous speed and accuracy, but only *that* task. LLMs are unable to go beyond their training or the prompt at hand, meaning they are tapping into a previously known set of information to respond to the user's prompt, despite the response being generated. Narrow AI models are unable to make new prompts without a user's input, or generate new knowledge that is not present in the databases available to the model (IBM, 2026). In spite of frequent updates and developments to the large language models available today, the crucial step of reasoning and deduction has not, or not yet, been achieved. But developers and tech sector leaders believe that, at some point in the near future, it will be achieved, creating Artificial General Intelligence (AGI). This is defined as AI capable of reasoning and deduction at a human level, capable of learning from previous tasks (IBM, 2026). This would render the conclusions of my argument at least questionable, but AGI is, at the time of writing, just a

theoretical construct. For this reason, analysis of the possible consequences of such a technological achievement is strictly beyond the scope of the dissertation.

2.2 AI and the Market in 2026

LLMs are a relatively recent technology, but their societal impacts are already being felt. The power and ease of use of LLMs have led to widespread adoption, as the OECD reports that over one third of its citizens used generative AI tools in 2025 (OECD, 2026), while a report by McKinsey & Company (2025) demonstrates that 88% of firms are using generative AI in some function. AI adoption shows that the technology has come to stay and is already causing societal shifts.

As was the case in the times of Marx and the Industrial Revolution, the issue of job replacement is central to the debate surrounding AI and LLMs. Current debate centres on the extent to which artificial intelligence will replace jobs or enhance human productivity in existing roles. As of June 2026, over 121 000 employees have been laid off worldwide in 2026 in the tech sector (*layoffs.fyi - tech and startup layoff tracker*, 2026), as firms increasingly replace their workforce with AI agents. This is not without controversy, as researchers suggest this phenomenon could also be one of necessary layoffs disguised under the cover of AI adoption in order to avoid reputational impacts (Caldwell, 2026).

Whether the impacts of AI on productivity are positive is still open to discussion. Academic research has indicated that the technology will reduce barriers to entry and increase productivity for low-skill jobs, while impacts on most high-skill jobs are still limited (Acemoglu, 2025; Brynjolfsson *et al.*, 2023). This suggests that AI could serve as an engine for small firms and individual entrepreneurs, simplifying tasks that previously required skilled labour. However, an International Monetary Fund article suggests that the increases in productivity created by artificial intelligence integration may have adverse effects on income inequality, if the wealth generated by increases in productivity is increasingly captured in capital markets (Rockall *et al.*, 2025). Whether productivity gains are large enough to offset capital retention will shape long-term impacts on inequality (Cazzaniga *et al.*, 2024).

3. From Mises to the Socialist Calculation Debate and beyond

Having established the technological concepts that I shall consider and how AI is already impacting the market, I now look back at the historical debate on markets vs. socialism and, specifically, their pricing mechanisms.

3.1 Main Austrian Arguments and their Opponents up to the late 1940s

In 1920, Ludwig von Mises published *Economic Calculation in the Socialist Commonwealth* (Mises, 1920), which presented socialism as an *impossible* proposition. Mises establishes the market's unique ability to transmit and concentrate dispersed, fallible information through prices, which are essential to informed decision-making. Those monetary prices result from the subjective preferences of individuals and signal whether they have, or have not, made an efficient use of their resources. Because socialism entails the abolition of private property – higher-order goods are *res extra commercium* – market transactions of the latter kind of goods cannot, by definition, exist. The absence of monetary prices for these goods entails that the socialist planners, however well intentioned, lack the signals they would need to make rational decisions.

The socialist reply came in the form of Oskar Lange's paper *On the Economic Theory of Socialism* (Lange, 1936, 1937), in which the author suggests a form of *market socialism* in which prices would be simulated by a committee of planners.

Prices for the list of basic resources – a list assumed to be given – would be accounting prices initially determined by the planners. State-owned firms would be told to take these prices as data and produce so as to minimise cost. Excess demand or supply of goods would certainly occur, but a trial-and-error process, in which planners would raise or lower the prices of their list of basic resources, would eventually lead to equilibrium prices.

Lange does recognise Mises' point that some form of accounting is needed to rationally allocate goods and services in an economy. But he distinguishes prices as "the exchange ratio of two commodities on a market", and as the "terms on which alternatives are offered" (Lange, 1936, p. 54). In his view, prices in this second sense suffice and are available under socialism, given information on production functions. He believes that the planners would have "exactly the same knowledge, or lack of knowledge, of the production functions as the capitalist entrepreneurs have" (Lange, 1936, p. 55). Having, moreover, data on preferences (a market in the institutional sense of the term would exist for consumer goods) as well as

the aforementioned list of resources, planners can simulate a market, correcting prices until equilibrium is reached, in a way inspired by Walrasian *tâtonnement* (though contradicting it in as much as there would be transactions before equilibrium prices were attained).

Lange's proposals were met with several rebuttals by Friedrich Hayek in the 1940s. Even earlier, Hayek (1937) rejects the pleonastic notion of 'given data' which Lange repeatedly uses. To whom are these alleged 'data' supposed to be 'given'? Do planners have access to a list of basic resources? Where is it? Markets, Hayek argues, allow the discovery of knowledge, dispersed across the economy's agents, and partly contradictory, rather than neatly presented in the planners' work sheet.

As Hayek explains in his 1940 article *Socialist Calculation: The Competitive "Solution"*, Lange's trial-and-error process does not make sense, as Lange does not go into the details of the process, such as the frequency with which prices would be updated by the central planners, which has implications on the firms' production plans. Hayek concludes that, in Lange's scheme, "the changes will occur later than they would if prices were determined by the market parties, and there will be less differentiation between prices of commodities according to differences of quality and the circumstances of time and place" (Hayek, 1940, pp. 134-135).

More importantly still, Hayek argues in *The Use of Knowledge in Society* (1945), Lange's argument reduces the calculation issue to one of computation – as explained above. In his words: "*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 problem which remains is purely one of logic" (Hayek, 1945, p. 519). But none of these "ifs" holds in the real world. It is through the market process that this information travels through the economy.

Hayek frequently uses the phrase "circumstances of time and space" (Hayek, 1945) to describe the type of local and specific knowledge he separates from more general scientific or technical knowledge. He uses the example of shortages of goods, such as tin: in case of a sudden scarcity or abundance of tin, agents with local knowledge of supply and demand will act accordingly, substituting the metal or increasing its use in order to adapt, in a process which will spread throughout the economy as more and more agents adapt. This happens, Hayek remarks, independently of the agents knowing why the supply or demand of tin has shifted, or indeed of the agents being aware of their actions' impact on the market. The sum of these actions is a form of communication through prices which spreads across the market

in a more efficient form, Hayek argues, than if a sudden market shift were communicated upwards to a central committee which would then make overarching decisions about the whole economy, by which time new facts would probably have affected the market, forcing the need for new analysis. To take this knowledge and its dissemination as given is to "assume the problem away", ignoring "the unavoidable imperfection of a man's knowledge and the consequent need for a process by which knowledge is constantly communicated and acquired" (Hayek, 1945, p. 530).

Hayek adds in 1948's *The Meaning of Competition,* that perfect competition and its static assumptions are irrelevant in an actual economy. Hayek sees the formation of knowledge as "a dynamic process whose essential characteristics are assumed away by the assumptions underlying static analysis" (Hayek, 1948, p. 94). Communication through the price system allows for knowledge to be shared, tested and revised, in a continuous process without agents consciously understanding they are doing so. The unknown nature of future knowledge yet to be discovered prevents its *ex ante* planning by any individual or institution.

In a socialist economy, Hayek argues:

> One of the most important forces which in a truly competitive economy brings about the reduction of costs to the minimum discoverable will be absent, namely, price competition […] the method which under given condition is the cheapest is a thing which has to be discovered, and to be discovered anew sometimes almost from day to day, by the entrepreneur, and that, in spite of the strong inducement, it is by no means regularly the established entrepreneur, the man in charge of the existing plant, who will discover what is the best method (Hayek, 1940, p. 139).

In short, Hayek sees market competition and the institutional environment behind it as the drivers behind the formation of new knowledge, directly contradicting Lange. Individual knowledge is subjective, contextual, partly tacit, and partly contradictory, therefore unavailable to a single mind or a central planning board. Without any planning or conscious action by the market's agents, this knowledge is disseminated in prices and partly corrected, i.e., in the market process, essentially a learning process, knowledge is discovered which otherwise could not even exist.

3.2. On recent, pre-AI Proposals of Socialism

Despite many failed attempts at socialism - none of which followed Lange, incidentally - the debate has remained alive, especially as new technologies, far beyond what was imaginable during the 1930s and 40s, emerged. Several proposals of socialism suggest using modern computing power and the internet to overcome Hayek's knowledge problem. Though Hayek's propositions go far beyond issues of data input and transmission, socialists have insisted that computers' processing power can solve the problem. One such case was Lange himself, in a 1967 paper published posthumously. At the time, computers were an academic reality, but still far from current standards. Lange defies Hayek: "Let us put the simultaneous equations on an electronic computer and we shall obtain the solution in less than a second. The market process with its cumbersome *tâtonnements* appears old-fashioned. Indeed, it may be considered as a computing device of the pre-electronic age" (Lange, 1967). It is apparent from these statements that Lange did not truly understand the essence of Hayek's critique. An overview of current proposals on the use of technology to overcome the knowledge problem is worthwhile to bridge the gap between discussions within the framework of traditional programming and in an AI context.

A proposal by someone who does understand Hayek is Evgeny Morozov's *Digital Socialism? The Calculation Debate in the Age of Big Data* (2019). Morozov proposes a form of socialism based not on central planning, but on a "digital feedback infrastructure" that is decentralised and automated, using big data. But Morozov's defence of Daniel Saros' (2014) *Information Technology and Socialist Construction* is difficult to comprehend. Saros' proposal is a complete flop, as Medeiros (2022) methodically demonstrates.

Jan Phillip Dapprich and William Paul Cockshott's article *Input-output planning and information* (Dapprich & Cockshott, 2023), as the title suggests, presents an input-output model of a centrally planned economy, where computation would perform millions of simultaneous equations. Dapprich and Cockshott argue that modern computers can process these data and plan an economy in seconds, since these data on inputs and outputs are already collected by most governments and firms. Despite not referring to AI in the model, their proposal is of interest to this argument due to the authors' stance on Hayek's "information argument" (Dapprich & Cockshott, 2023, p.419). The use of the word *information* as opposed to *knowledge* hints at why the authors, much like Lange decades before, do not directly reply to Hayek.

Despite agreeing that collecting comprehensive data on the production of millions of different products would be "a significant undertaking" (Dapprich & Cockshott, 2023, p. 419), the authors believe that this should be possible with modern technology, at least in principle. Dapprich and Cockshott consider that this information "can be represented in a table and is thus not tacit" (Dapprich & Cockshott, 2023, p. 419). This information would then serve as a basis to plan the economy, which would evolve alongside technology, with forecasts on technical evolution left to experts. As processing power evolves, the plans become more agile, as the planners would calculate necessary changes to the plans faster. In a two-step approach reminiscent of Lange's trial and error approach, consumer goods would be sold with prices being adjusted until supply matches demand. Then, these prices would be compared to labour values, using the input-output tables. Production levels are adjusted until labour values are equal to prices. Recognising (like Mises) that labour values may be inadequate to measure production costs, Dapprich and Cockshott (2023, p. 421) suggest the use of shadow prices. Powerful computing and the evolution of the planning mechanism would allow for the knowledge problem to be surpassed.

The authors do recognise that some knowledge, which they call "how-to knowledge" (Dapprich & Cockshott, 2023, p. 420), is tacit. But this knowledge does not need to be centralised, as such details of production should be irrelevant to the resource allocation made by the planners. The authors surmise that, regarding the processing of consumer data, this process is informationally equivalent to the market process.

Some of Dapprich and Cockshott's assumptions are clearly quite controversial from an Austrian viewpoint. It is of little surprise, therefore, that this paper sparked a discussion, namely in Lambert *et al.* (2023). Since Dapprich and Cockshott do not directly tackle the issue of AI, an analysis of these criticisms is beyond the scope of this dissertation.

3.3 New Proposals: How could AI replace the Market?

Now, we turn the page to AI. Some authors have already presented some proposals worth noting. At this stage of the discussion, debate comes from both the economics and the technology areas.

Coming from the technology field, Spyridon Samothrakis wrote two articles on how AI can be used to plan an economy. *Artificial intelligence inspired methods for the allocation of common goods and services* (Samothrakis, 2021) is his first foray into the topic. Interestingly, the article was

written before the AI boom truly started. In it, Samothrakis proposes the automated planning of parts of an economy through artificial intelligence, with an emphasis on the technical feasibility of the model. Believing markets to be flawed and inefficient for the production of certain essential products, Samothrakis suggests a model, based on input-output tables, which he names "Open Loop In Natura Economic Planning - OLIN-EP" (Samothrakis, 2021, p. 5).

OLIN-EP operates based on a reinforcement learning mechanism, whereby an AI agent identifies which state it is operating in, by capturing data on the economy transmitted by producers and consumers. Afterwards, the agent takes actions to modify the state, based on its programming and the power it is given to plan the economy. The objective of the AI agent is to maximise its reward function, which is also programmed into the model. How the economy reacts to the agent's actions is captured in the transition function, which is unknown in this simulation.

The author seeks to distance himself from traditional socialist literature, reinforcing that this model is not an attempt to replace the market entirely, but rather to distribute essential products directly to consumers. The OLIN-EP should plan the economy so that demand for these products is met, on a global scale, with daily revisions incorporating the previous day's data on production and demand. Like the models above, this depends on the availability of said data. Production data would be provided by the production units, who would fill the input-output matrix assisted by the model's machine learning algorithm. Optimisation of this process is expected by Samothrakis, but not thoroughly explained.

As for consumer data, the AI agent would predict future consumer preferences by learning from past and present transactions. The agent's reward function would be based on past transactions, and future learning would occur through the agent's reinforcement learning mechanism. Additionally, direct questions and voting should assist in this process. Samothrakis (2021, p. 10) states that "this should allow for effectively the discovery of basic needs on a fundamental level and the provision of relevant goods". Tailoring to individual consumers is technologically inviable, while insufficient personalisation of goods will lead to complaints. Therefore, Samothrakis suggests that "the focus should be on goods that allow for a high degree of customisation" (Samothrakis, 2021, p. 10). Production of customisable goods would solve the issue of lack of variety, as consumers could customise the goods they

purchase. The development of new technologies allowing for customisation of products is assumed.

Since the market will continue to operate for all non-essential goods, Samothrakis highlights the need to minimise abuses and black markets, as well as using the market to produce improvements to the planning mechanism. Importantly, markets should be made to guarantee that productivity gains of this plan are geared towards increasing wages. The model's simulations show that it is technically very simple to plan up to 50 000 industrial goods and 5 000 final goods in a personal computer.

Overall, the model is, as the author plainly states, simplistic. It is meant to be a first step towards more focussed research, especially as artificial intelligence technology develops. Among the limitations acknowledged by Samothrakis, the one regarding "sector heuristics" deserves to be highlighted. Samothrakis states that research is needed in order to discover heuristics that shape economic agents' behaviour in individual sectors. This is implicitly an acceptance that the model proposed fails to discover knowledge as effectively as the market. In the market system, these heuristics are tacitly discovered through interaction among agents.

Despite recognising that Hayek's "epistemic argument […] is still very valid" (Samothrakis, 2021 p. 4), Samothrakis believes Hayek's view fails to acknowledge that a sizeable part of the population has similar preferences for essential goods, such as food, shelter, among others. This line of reasoning is flawed, though. Some items listed as essential, such as basic communication devices or electricity, are considered essential now but surely not before they existed. The innovations that eventually become essential to basic human consumption are discovered through the market process, as their invention and use cannot be planned *ex ante*.

In 2024, Samothrakis published a new article, *Artificial intelligence and modern planned economies: a discussion on methods and institutions* (Samothrakis, 2024). Here, he shifts from discussing a model, as in his 2021 paper, to discussing his vision for the institutions of centralised economic planning. Samothrakis recognises that despite increased societal interest, AI is still far away from replacing human labour, primarily due to its inability to replicate human reasoning and the number of marginal events that elude its algorithmic predictions. Furthermore, the author sees the prospect of a revolution such as the abolition of the market as unlikely, so he opts to chart a path of small changes in order to strive towards a more egalitarian distribution of wealth in society, using AI as a tool.

To this end, he once again turns to a reinforcement learning (RL) mechanism. The dynamics of this mechanism are similar to the previous model:

> An agent assesses the world through its senses and can infer in what state $s \in S$ it is in. It can move to a new state using actions $a \in A$. A transition function $P^a_{s,s'}$ decides in which state s' an agent will move into after taking action a. When an agent performs an action, it receives a scalar reward R^a_s, which denotes the amount of "pleasure" an agent receives by performing this action in the specified state. Finally, a discount factor denotes how the agent treats rewards time-wise (Samothrakis, 2024, p. 2963).

Samothrakis incorporates Marxist concepts into this dynamic: AI agents would invent and produce commodities with the objective of reward maximisation. The reward is the value of the commodity, measured in either monetary or use values. In this context, he discusses the labour theory of value: if the AI agent captures all the reward, it is stealing since all value is derived from labour and nature, citing Marx (1875). However, the author accepts that this view is far-fetched: "Labour alone without a concept is obviously just a waste of time and effort" (Samothrakis, 2024, p. 2964).

Samothrakis suggests three main goals for the use of AI to improve central planning: "commodity power", "decentralisation of production", and "equivalence of process" (Samothrakis, 2024, p. 2967-2968). *Commodity power* refers to the idea that the economy should create what the author describes as "empowering commodities" (Samothrakis, 2024 p. 2967) – commodities that can be used to create several other commodities, creating decentralised hubs of local production where individuals are empowered to produce these commodities. *Decentralisation of production* is connected to this. The planning institution should have the primary role of coordinating these hubs, as opposed to commanding as a traditional planning committee would. Finally, *equivalence of process* describes a focus on having these decentralised production hubs be as similar as possible, attempting to create harmonious final commodities. Samothrakis explains:

> The more commodities are moved into the "equivalent of process" bucket, the less one would have to discard during production and the easier the production process becomes, as generally we would not care that much about variance in the outcomes (e.g. you would probably end up making your coffee slightly different each morning if the machine you used to make it was self-assembled; even if we assume that the

same instructions were followed on how to make dress, home-made ones would still not look exactly the same) (Samothrakis, 2024, p. 2968).

To achieve these goals, the institution in charge of planning would work in five areas: value identification (R&D), value protection (ensuring the supply of basic necessities), value mixing (maximising intertemporal value), value ethics (oversight) and meta-value (evaluating the programme).

Samothrakis' vision entails a network of small producers, who own the means of production they are using but work under the coordination of an automated AI system. In keeping with the previous article, this network would handle the production of products taken as essential, which raises the same questions as before: who determines what products are essential? The article presents a way for AI to assist in some coordination of a complex economy, by using reinforcement learning to recognise patterns in a complex economy and increase coordination between different producers, but it says very little on how an AI model could centralise price-setting decisions in an efficient way.

3.4. Austrian Proposals on AI

Against these proposals, there is already some literature in the field of Economics opposing the idea that AI will make centralised price-setting possible. In 2023, a special edition of the *Journal of Economic Behavior and Organization* (Phelan & Wenzel, 2023b), opened the discussion, featuring several contributions to the topic, such as Phelan and Wenzel (2023a), Lambert and Fegley (2023) and Boettke and Candela (2023). This edition also includes the debate on Dapprich and Cockshott (2023) I mentioned above. Further contributions came from Davidson (2024, 2025), Gmeiner and Harper (2024), and Makridis and Lazanski (2025). I will highlight some of the arguments presented.

Both Lambert and Fegley and Boettke and Candela restate the Misesian argument on economic calculation, arguing that AI and other modern technological developments are unable to solve the problem Mises identifies. But they focus on proposals for 'cyber-communism' and 'technosocialism', respectively, which do not take the arguments set forth by Mises and Hayek seriously.

Phelan and Wenzel outline some of the issues at hand - such as innovation, attention and judgement - when discussing the role of markets in the context of recent technological advancements (big data, quantum computing and AI), highlighting AI's lack of true agency.

An AI model cannot efficiently innovate due to its decisions being based on its programming, which comprises past information. This information is vast and the algorithm behind the model would have to be able to decide what part of it is relevant and make judgements on what to do when given the data. Both actions require the model to be programmed to act in certain ways, raising the question of *who* programs the model. In effect, the programmers behind the model would be *de facto* planners, as the priorities they provide the model with would dictate how it behaves.

Gmeiner and Harper perform an adapted Turing test (Turing, 1950) on AI. The Turing test is a benchmark to test whether a *machine* can *think* like a human. The authors reframe this to check if AI can plan an economy and compare it to the free market economy. Firstly, they state the basic requirements for the experiment to work: the model must be fed with reliable and processed information, with minimal discrepancies between data sources. This is a major obstacle, as the process of cleaning up raw data requires human and technological labour. Then, Gmeiner and Harper tackle the decision-making aspect of the model, where they identify several requirements. The model must be able to interpret its state space, ergo the world it lives in, and the current state of the economy measured up against the model's objectives, as well as its control space, comprising the parameters it can change in order to control the economy. It must be able to take into account the rest of the world, as its economy will interact with other economies. The time horizon at hand is also a key variable, as it will shape the model's decisions. The model must be able to identify and correct inefficiencies like the market process does and, finally, it must have the power, granted by political power, to act upon its determinations. Ergo, it must be empowered to plan the economy as it sees fit. In this model, there must be a hierarchy of goals, so that the model can optimise which goals to achieve and how to achieve them, in contrast to the spontaneous order of the market.

The results of this model, the authors argue, create several public choice problems. The main issue is, again, *who* controls the model. The model's goals must be set by planners, who would gain tremendous power in this paradigm. The planners would be elected officials, setting goals that programmers would configure into the model. Therefore, programmers would also have enormous power, as they have the effective control over the model despite the officials having political responsibility. Programming being beyond the expertise of planners, the former will have a considerable degree of independence from political power. Potentially,

they could manipulate the model for personal gain. A second public choice problem arises from data sources, which are controlled by other entities - private firms in charge of networks, software, among others. There is an incentive for these firms to distort the data supplied to the model to their advantage. In any event, the model would be unable to enact judgement on which outcome is optimal in the absence of market prices.

Sinclair Davidson has written two articles on the topic, taking an Austrian perspective to review if AI could organise an economy as efficiently as the market. In 2024, he published *The economic institutions of artificial intelligence* (Davidson, 2024), tackling the issue from a perspective of bounded rationality. Davidson argues that artificial intelligence may solve issues related to the heuristics and mental shortcuts that are intrinsic parts of human decision-making due to its established ability to rapidly process vast amounts of information. However, the author separates information from knowledge, as used by Hayek. AI models are apt at processing information, but unable to create knowledge beyond their initial programming and therefore cannot participate in the process of discovering new knowledge, an essential product of the market process. Additionally, they are unable to absorb tacit knowledge, which must be transformed, at a cost, into information that the model can learn. Therefore, AI cannot centrally plan an economy.

A year later, in 2025, Davidson published *The limits of artificial intelligence* (Davidson, 2025), diving into the topic of AI again to discuss the function of the entrepreneur in the context of artificial intelligence. After restating AI's inability to surpass the knowledge problem, Davidson moves on to discuss creativity and how AI-generated content, despite looking like an original creation, is actually a recombination of existing content, reprogrammed into a different form. This stands in contrast to human creativity, which is intentional and much deeper. Creativity, Davidson claims citing Bylund (2025), is a core part of the entrepreneurial function and the key reason he uses to argue that AI is unable to fulfil it. The author then presents the Theory of the Adjacent Possible by Koppl *et al.* (2023), which states that technological evolution occurs not as a linear process, but as a process of recombination of existing technology. This would *prima facie* present an opportunity for AI to fulfil the entrepreneurial role, as its essential function is the recombination of information in its programming. However, even if one accepts this theory, an AI model would still have to enact judgement in selecting which combinations are worthy of being chosen. Since the

context required to perform these judgements, such as tacit and local knowledge, is not programmable *ex ante* into the model, entrepreneurship remains an impossibility.

In *Navigating the Knowledge Problem in the Era of Open AI and Polycentric Governance* (Makridis & Lazanski, 2025), Christos Makridis and Dominique Lazanski tackle problems surrounding the adoption of AI and discuss forms of AI governance. The authors restate that artificial intelligence, due to its preset programming, cannot capture tacit knowledge or produce new knowledge in a dynamic setting, citing Boettke and Candela (2023) and Davidson (2024). In their view, "praxeology […] reinforces Hayek's assertion that centralised governance structures - whether in planning entire economies or designing AI oversight - are fundamentally at odds with the dynamic, subjective nature of human action" (Makridis & Lazanski, 2025, p. 9).

Makridis and Lazanski retrace Boettke, Candela and Davidson's arguments. An interesting contribution to the cited works is the discussion of property rights in an AI context. LLMs, despite being prompted to *generate* content, can only produce content that is derived from existing data, which is harvested from content available to the public. This process occurs without any compensation for its creators, creating a hitherto unsolved issue of copyright and content ownership, as current intellectual property rights do not take AI content generation into account. The ownership of the generated content becomes unclear, and the decentralised processes of innovation and creativity come into question as the process of content *creation* by individuals is replaced by *generation* through LLMs. As the authors explain, "this dynamic reflects a broader governance failure: LLMs are currently treated as commons without well-defined ownership structures, eroding the market signals that normally guide the production and valuation of intellectual labour" (Makridis & Lazanski, 2025, p. 3).

The authors identify this as an AI governance issue, which they use to highlight the need for proper AI regulation. To this end, a framework of polycentric governance is proposed. Under this framework, AI would operate with several separate decision centres, rather than one centre or with one set of rules. Developed by Vincent and Elinor Ostrom (Ostrom, 1990), polycentric governance operates under five principles: (1) 'clearly defined boundaries', (2) 'collective-choice arrangements', (3) 'monitoring and sanctions', (4) 'conflict-resolution mechanisms' and (5) 'nested enterprises'. To summarise, the framework is based on local collaboration to form decision centres with well-defined limits, objectives set by local stakeholders, strong oversight and arbitration mechanisms, all operating under a wider

institutional umbrella handling high-level issues that affect several sectors or centres. Several elements of this structure are already operating in practice, so the authors believe that this framework will be implemented as part of the natural progression of AI evolution.

Makridis and Lazanski identify this framework as a roadmap for AI governance in a way that can preserve the integrity of localised knowledge. Opposing central planning institutions, the authors "argue for institutional arrangements that preserve the option value of adaptation: rather than enforcing top-down uniformity, effective governance should provide space for bottom-up discovery, customisation, and iterative learning" (Makridis & Lazanski, 2025, p. 12).

But the tacit nature of knowledge and the dynamic and spontaneous nature of human adaption to new knowledge remain obstacles current AI technologies cannot surpass. Therefore, the framework Makridis and Lazanski describe serves to facilitate AI integration as an aide to current market mechanisms, as opposed to replacing them.

4. On the Feasibility of AI Replacement of Markets

From bartering economies to e-commerce, the market has been a staple of human societies for centuries. The market economies it has powered have flourished into advanced countries with high living standards and increasing wealth. Transactions may have evolved in their complexity and context, but the core elements remain unchanged: economic agents trade goods and services amongst each other, in exchange for a means of payment. Technology may evolve at a rapid rate, but so far it has been unable to replace the market.

Some of the proposals I analysed above ponder the possibility of planning an economy with some form of algorithm or LLM. The existing body of work presents interesting points, but no conclusive answer on how markets will be impacted by the growing phenomenon of artificial intelligence. It is in this unstable context that I will try to assess whether artificial intelligence could replace the market.

Should AI replace the market, transactions of goods and services would be centralised within an AI model, presumably an LLM or agent. Let us assume that it is given sufficient processing power to plan the economy in real time, which is no small task. Immediately, two questions arise: who oversees the programming of the model, hence who controls it - private firms or the government - and what would its programmed objectives be? Both these public choice problems are discussed in Gmeiner and Harper (2024).

4.1. Speculations on the End of Markets

The process of substituting the market could be a voluntary one, whereby agents freely surrender their production and consumption decisions to the AI model, as opposed to trading on the market, or it could be imposed either by the state or by powerful firms. I will look at these aspects in turn.

Let us assume an LLM is developed – for now, the question of who develops it is not of importance – that can plan an economy and process its transactions. A voluntary substitution of the market would take place if agents willingly perform transactions through this LLM instead of the market. For this transition to occur, the LLM would have to be more efficient than the market in transmitting agents' intentions, higher efficiency meaning that their needs would be better fulfilled in this new system.

Alternatively, the shift to an LLM could be imposed by the state or by big firms. If it were imposed by the state, the shift would occur through legislation, presumably in a context of wider societal revolution, as the state would have to seize control of the means of production to control the economy.

A voluntary substitution is unlikely, as it would take a collective decision by societies worldwide to relinquish their power over consumption and production. A replacement through state power could only occur in a context where liberal democracy collapses, to be replaced by a *de facto* dictatorial regime where power over economic decisions is mostly centralised. However, it is far-fetched to believe that this could occur at a global level and, as history has proved repeatedly, such experiments – the socialist experiments that history has witnessed – have failed.

A substitution through corporate power is also conceivable (and likewise a threat to democracy). I find this to be the likeliest of the three replacement scenarios. In such a scenario, private ownership of the means of production is maintained, but transactions are made through an LLM. Of course, such a change entails further implications regarding how transactions would take place, which I will address later on.

The current market landscape is dominated by large tech firms, such as Microsoft, Meta or Alphabet. These large firms control large shares of their markets on a worldwide scale and have often turned to power consolidation through horizontal acquisitions (Motta & Peitz, 2021). The market value of these firms is, in part, dependent on their investment into AI research, so they are very interested in the successful implementation of artificial intelligence in society. As of July 2026, the two most advanced artificial intelligence firms, OpenAI and Anthropic, are valued at approximately $1 trillion each (Bergman, 2026). Mutual investment among tech firms such as these is frequent, generating trillions of dollars in deals (Petrova, 2025). It stands to reason that one or several firms would design and control the LLM that would replace the market in the scenario I am considering. The combination of financial and technological power that these firms wield can be used to force people to participate in the replacement of the market. The cost of non-participation could be exclusion from these platforms, which are crucial in modern life. Of course, this is quite different from enforceable law, but the control these firms hold over digital infrastructure and the importance of the digital world in modern society mean that the tech firms in question wield considerable

power. Moreover, large firms would likely use their influence to sway political power into forcing such a change through legal means.

Authors like Durand (2020) and Varoufakis (2023) have coined the term *technofeudalism* to describe the increase in power held by tech firms in recent years. In the past, feudal lords controlled land, renting it or allowing peasants to live and work on it. Nowadays, tech firms extract tremendous amounts of data from consumers and control the digital platforms through which millions interact daily. These data are then sold, mainly for marketing purposes, forming an important revenue stream for some of the wealthiest firms in the world and raising controversial questions on data ownership and customer privacy. Whether the term applies to current market dynamics is beside the point of this dissertation, but I believe that some variant of *technofeudalism* is a likely outcome in a world where an LLM centrally controls the economy. The firm or firms controlling the model would become *de facto* rulers, overshadowing any government.

There is also the possibility that the LLM, through its unsupervised learning methods, learns to disobey whoever controls it. This would be a question normally reserved for science fiction, but it is nevertheless most interesting. If the power to make decisions on production, investment and consumption is given to an AI model, it is fair to ask if humans are in control of the model or if (and when) the model can wrestle control away from humanity.

But substitution of the market is not the sole possibility: AI could continue coexisting with the market. Perhaps, like other technologies before it, it would not have succeeded in replacing the market. This scenario may be the most likely, with firms continuing to use AI to attempt to influence market dynamics. Even the likeliest form of market substitution, through corporate power, relies on the (perhaps still somewhat implausible) premise that existing firms are powerful enough to enforce such a change.

4.2. Price-Setting under AI

The issue of prices is central to any debate on the role of the market. Any convincing proposal to replace the market must include a comprehensive explanation on whether prices exist in such a conception, what is their nature and how they are formed. In this section, I look at these questions, discussing whether they may present obstacles to the substitution of the market through artificial intelligence.

The market process has continually evolved, even if it is hardly flawless. Nowadays, technology allows pricing decisions to be personalised, *i.e.*, adapted to the individual consumer, in certain transactions – especially in online environments. Let us take the example of Uber, the largest rideshare app in the world, which uses dynamic pricing: a ride can have different prices for different consumers, depending on factors such as the time of day, the consumer profile and even the battery level of the consumer's device (Vedantam & Penman, 2016). Such a case of widespread price discrimination by a single provider is a clear example of recent technologies influencing market processes.

But the proposals I analysed in Chapter 3.3. are not clear as to how an AI system could produce rational prices. Samothrakis does not focus on the economic details of his proposals, opting instead to provide his vision of a society in which AI manages the production of *essential goods*. In suggesting a system where AI planning and the market coexist, Samothrakis wishes to build a bridge to current market societies: since he sees replacing the market as unlikely, a compromise is proposed. However, the author fails to explain what *essential goods* are and who determines their importance to society. Additionally, no detail is given as to how AI and the market are supposed to interact when the same resources are required to produce both market and non-market products. In the absence of a market price offering producers a rational way to rank offers, no obvious way to allocate disputed resources exists. The allocation can be made by the AI model, but this would render the entire economy centrally planned. Alternatively, if market prices for factors of production exist, then the production hubs Samothrakis proposes would be producers operating in a market, rendering the planning mechanism obsolete.

Since a system featuring AI planning and a market simultaneously would apparently not work, I will now focus on a scenario of full substitution of the market. Firstly, it is important to define what a full substitution of the market entails, with a special focus on the issues of private property rights over means of production. The substitution of the market must entail an abolition of those private property rights, which would create a society closely resembling socialist proposals such as those I discussed earlier, with the AI model or agent as the central planner. Without private property, means of production do not have prices and are controlled by the state. Let us assume they are controlled by the LLM in charge of planning the economy. Without market prices for these factors, the model cannot efficiently compute the costs of the decisions it is taking, which renders those decisions useless. Mises' argument

holds over one hundred years later: rational economic decisions cannot happen without market prices, a problem that no artificial intelligence mechanism hitherto invented can solve.

The scenario in which factors of production are privately owned is far more interesting. Firstly, if these factors are privately owned, they must be acquired from their owners, hence there must exist prices. Lange's market socialism used accounting prices set by the planners. In a similar vein, the AI model could fix prices, correcting them afterwards as it deems necessary. The immediate advantage of this conception in comparison to market socialism is the speed of the trial-and-error process of price correction. An AI algorithm can adjust prices in real time, since it can accumulate and process large amounts of data. As Dapprich and Cockshott (2023) suggest, these data are already collected in modern transactions, especially in online settings. A centralised artificial intelligence would harvest data on all transactions, creating a real-time picture of the economy that would be unimaginable in Lange's time. The use of these data should accelerate revisions of prices. Adjustments to prices would be dictated by the AI's objectives. However, the use of simulated prices yields several problems, analogous to those Mises and Hayek identified.

Firstly, questions of incentives arise. Producers are entirely dependent on the AI planner's decisions. Why should they favour an environment where profits depend on an AI planner? Moreover, considering wages are the price of labour, a plethora of questions arises when considering the implications of set prices in the labour market. Most notably, how will productivity be incentivised? This would depend heavily on the model's programming and objectives, especially if the latter are set by political power. Furthermore, to assume that the computing power exists for an LLM to calculate wages down to the individual worker is questionable.

Moreover, there is the risk of the model's reinforcement learning mechanism leading it to bad planning practices. Since reinforcement learning occurs when the model learns based on its own past actions, inefficient planning could become self-sustaining if the learning mechanism does not identify it as a behaviour to eliminate. This plays into the power dynamics explained above, raising the question of whether humans are still in control of the model, or vice-versa.

Most importantly, prices determined by an LLM are as useless as prices determined by a central committee, since they lack proper rationality and ranking in deciding between

possible scenarios for the economy. The use of AI instead of a committee does not invalidate Mises' argument that the absence of prices is equivalent to abandoning rationality.

Looking back at Hayek's critique of Lange: "*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 problem which remains is purely one of logic" (Hayek, 1945, p. 519). These *ifs* Hayek enumerates remain unknown. Even if the planner is a powerful AI algorithm, it is not omniscient. Let us take, for example, the information on how many trees remain in the Amazon rainforest. No market or planner will know the exact quantity of trees there are at a given time, but while the market mechanism can continuously adapt the price of wood through the interaction of local lumberjacks and salesmen with the wider market, an AI planner cannot know local conditions unless they are explicitly programmed. It is impossible for the model to know, or even discover, economic agents' preferences due to the nature of knowledge, which I will now discuss.

4.3. The Nature of Knowledge

Knowledge is the result of a dynamic process of discovery, the market process. Returning to the example of the Amazon rainforest, as Hayek wrote, it is the "man on the spot" (Hayek, 1945, p. 524) that best knows the specificities of local conditions that may come to impact the price of wood extracted from the rainforest. He has this largely tacit knowledge, accumulated through experience and interaction with local conditions. This knowledge, a kind of 'guess', may be entirely contradictory with 'guesses' of others. But interaction with the market will validate or falsify these 'guesses' or fragments of knowledge. Knowledge will be discovered or revised and disseminated. A local factor such as the weather will impact the prices charged in the early production phases, trickling down to the final consumer not only of wood, but also of derivatives that use wood in their production. This transmission happens seamlessly and without any need for conscious action, through market prices whose movements communicate to the parties in a transaction that a change has occurred somewhere along the value chain. As transactions occur, new knowledge is constantly discovered, communicated, tested. This dynamic process is unpredictable: it is impossible to know in advance what knowledge will be discovered at any given time. What is taking place is a learning process which has not been designed by any human mind.

From this framework, a framework that so far no one has managed to falsify, it follows that AI cannot really *discover* knowledge. In Chapter 2, I observed that modern artificial

intelligence is advanced, but still a form of *narrow* AI: it cannot expand beyond its training and the prompts given to it. An LLM can process existing information with tremendous efficiency, but it cannot discover new information that is not explicitly programmed or prompted into it. This is not to say that an LLM cannot learn information without new programming, but it cannot ascertain whether a piece of information is relevant (or not) to a task. The model works, as I explained earlier, in a probabilistic manner which may fool the user, but does not constitute reasoning. An LLM can easily scour news websites for the most recent news worldwide, therefore integrating new information into its database. It can produce an accurate guess, based on its training, that a piece of information is relevant to its task, but it cannot decide by itself what information is relevant without a user's input or previous programming. Such determinations are made tacitly by humans, as part of our reasoning and decision-making skills. These abilities are beyond the reach of artificial intelligence. To summarize, AI cannot *think* at a human level.

The distinction in how humans and AI *think* is the key to this argument. Knowledge is local and mostly tacit. Frequently, humans discover knowledge without realising it, through their interactions in the market. Price signals allow for a form of unspoken communication through which agents who may never meet in any way understand one another, connecting the aforementioned "man on the spot" with the final consumer. This is a connection that AI cannot spontaneously make.

Firstly, an LLM is not *local*. The model can be taught, by a user, that an event occurred, or it can find a news article reporting it (though it cannot make a decision on whether the source is reliable other than, again, a probabilistic guess), if an event is newsworthy. But it cannot directly observe an event and react to it. In fact, it may even be entirely unaware that an event occurred, as it is impossible for the LLM to be omniscient. 'Knowledge' can be discovered at any time by a single person and, in the absence of the market mechanism, it cannot be validated or falsified.

Furthermore, an LLM cannot integrate tacit knowledge without it being expressly communicated to it, which presents a contradiction. Tacit knowledge is, by definition, implied. It can be conveyed in written form - just as explicit knowledge can become tacit in time - but this written form may not be very useful: as my supervisor used to remark in lectures, do we learn to cycle by reading a text explaining how to do it? The market process incorporates tacit knowledge seamlessly, an LLM does not.

Returning to the example of the Amazon Rainforest, an LLM can make the correct guess that an increase in wildfires in a given year will impact the price of wood. However, it first needs to learn that the number of wildfires has increased, either through direct communication by a user or through news outlets. Afterwards, it scours its database and correlates previous wildfires and price movements. This is the part of the process that LLMs excel at. After correlating these data, it can produce a determination based on them which maximises its reward function. In practice, this is a standard case of backward-looking behaviour. What an LLM cannot replicate is the forward-looking behaviour that agents, especially entrepreneurs, display in their interactions with the market. It lacks the ability to make the tacit connections that a human mind makes when presented with such information: in the face of price signals, an entrepreneur or a producer may spot a profit opportunity through new products or methods, taking a risk with the intent of capturing this profit. An LLM, however, cannot conceive new ideas beyond its programming, only recombine existing concepts. It can only make a guess, based on previous data points, that an event may have repercussions on a price and provide a reward-maximising outcome. Since that guess is based on *ex ante* learning, the LLM cannot react to new events or create new products. Even *generative* AI, as Makridis (2025) highlights, cannot *generate* new content, but in fact only recombine existing content. As Davidson (2025) illustrates, recombination is neither innovation nor entrepreneurship, as it lacks the creativity and intent of human innovation.

To summarise, an AI model does not *think* as humans do. Strictly, it does not think at all. It takes a probabilistic approach that yields a result that may resemble thought but is an entirely different process. Therefore, AI in its modern form cannot solve the knowledge problem. The inability to discover knowledge presents an unassailable epistemic barrier that narrow AI cannot surpass.

As I see it, then, AI is not a viable replacement for the market. The prospect of substitution of the market is one that most likely fits an authoritarian future in the hands of an all too powerful state or of oppressive corporations, and devoid of economic rationality. It is unlikely that such a society would not replicate too many vices of the socialist experiments of the 20th century.

5. Conclusion

In 1979, IBM published a handbook which contained a now-famous quote: "A computer can never be held accountable, therefore a computer must never make a management decision" (Bonderud, 2025). At the time, computers were primitive and limited in their abilities. Forty-seven years later, artificial intelligence is at the forefront of technology, but it is precisely that: *artificial*. It still lags behind the human abilities to think and reason, fooling its users with a statistical approach that attempts to imitate human thinking. AI could theoretically be programmed to manage a firm, but it still cannot make decisions of its own or face their consequences. This is why the arguments set forth by Mises and Hayek remain relevant, over one hundred years after Mises first published *Economic Calculation in the Socialist Commonwealth*: the market remains the best mechanism to disseminate and validate knowledge and promote economic rationality.

The Austrian arguments present the most powerful defence of markets, but the nature of a seismic shift such as the substitution of markets invites studies from several points of view. Of course, as the technology evolves, research into future forms of AI and the market remains essential.

After researching and writing this dissertation, I firmly believe that a future where markets are replaced by AI is bleak. It would, certainly, be authoritarian and a far cry from the modern liberal democracies that the market supports. Barring the development of general AI capable of reasoning, it would certainly be riddled with the same inefficiencies as socialist experiments of the 20th century.

6. Bibliography

Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, *40*(121), 13-58. https://doi.org/10.1093/epolic/eiae042

Bergman, B. (2026). Anthropic has soared to a $1.2 trillion valuation on secondary markets. Shares are impossible to get. *Business Insider*. https://www.businessinsider.com/anthropic-soared-1-2-trillion-valuation-on-secondary-markets-2026-7

Bergmann, D. (2025). What is Machine Learning? https://www.ibm.com/think/topics/machine-learning

Boettke, P. J., & Candela, R. A. (2023). On the Feasibility of Technosocialism. *Journal of Economic Behavior & Organization*, *205*, 44-54. https://doi.org/10.1016/j.jebo.2022.10.046

Bonderud, D. (2025). AI decision-making: Where do businesses draw the line? | IBM. https://www.ibm.com/think/insights/ai-decision-making-where-do-businesses-draw-the-line

Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative ai at work. In: National Bureau of Economic Research.

Bylund, P. L. (2025). *Entrepreneurship and evolutionary economics*. Cambridge University Press.

Caldwell, S. (2026). 'AI is now the leading reason companies give for cutting jobs,' says new report—what that means for workers. *CNBC*. https://www.cnbc.com/2026/06/05/ai-is-now-the-leading-reason-companies-give-for-cutting-jobs-says-new-report-what-that-means-for-workers.html

Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., & Rockall, E. (2024). *Gen-ai: artificial intelligence and the future of work*. International Monetary Fund.

Copeland, B. J. (2026). Artificial intelligence (AI) | Definition, Examples, Types, Applications, Companies, & Facts | Britannica. In. https://www.britannica.com: Encyclopædia Britannica, Inc.

Dapprich, J. P., & Cockshott, W. P. (2023). Input-Output Planning and Information. *Journal of Economic Behavior & Organization*, *205*, 412-422. https://doi.org/10.1016/j.jebo.2022.10.041

Davidson, S. (2024). The Economic Institutions of Artificial Intelligence. *Journal of Institutional Economics*, *20*(1), 1-16. https://doi.org/10.1017/S1744137423000395

Davidson, S. (2025). The Limits of Artificial Intelligence. *The Review of Austrian Economics*. https://doi.org/10.1007/s11138-025-00705-2

Durand, C. (2020). *Techno-féodalisme: Critique de l'économie numérique*. Zones.

Gmeiner, R., & Harper, M. (2024). Artificial Intelligence and Economic Calculation. *AI & Society*, *39*(3), 985-1007. https://doi.org/10.1007/s00146-022-01616-3

Hayek, F. A. (1937). Economics and knowledge. *Economica*, *4*(13), 33. https://doi.org/10.2307/2548786

Hayek, F. A. (1940). Socialist Calculation: The Competitive `Solution'. *Economica*, *7*(26), 125-149. https://doi.org/10.2307/2548692

Hayek, F. A. (1945). The Use of Knowledge in Society. *The American Economic Review*, *35*(4), 519-530. http://www.jstor.org/stable/1809376

Hayek, F. A. (1968). Competition as a discovery procedure. In *The Quarterly Journal of Austrian Economics* (Vol. 5, pp. 9-23). https://cdn.mises.org/qjae5_3_3.pdf

IBM. (2026). Understanding the different types of artificial intelligence. https://www.ibm.com/think/topics/artificial-intelligence-types

Kariuki, N. (2025). *Artificial Intelligence Index Report 2025*. https://hai.stanford.edu/ai-index/2025-ai-index-report/economy

Koppl, R., Gatti, R. C., Devereaux, A., Fath, B. D., Herriot, J., Hordijk, W., Kauffman, S., Ulanowicz, R. E., & Valverde, S. (2023). *Explaining technology*. Cambridge University Press.

Lambert, K. J., & Fegley, T. (2023). Economic Calculation in Light of Advances in Big Data and Artificial Intelligence. *Journal of Economic Behavior & Organization*, *206*, 243-250. https://doi.org/10.1016/j.jebo.2022.12.009

Lambert, K. J., Fegley, T., Candela, R., Boettke, P. J., Phelan, S. E., Wenzel, N. G., & Dapprich, J. P. (2023). Reply and Counter-Reply: On Cybersocialism. *Journal of Economic Behavior & Organization*, *212*, 300-310. https://doi.org/10.1016/j.jebo.2023.03.026

Lange, O. (1936). On the Economic Theory of Socialism: Part One. *The Review of Economic Studies*, *4*(1), 53-71.

Lange, O. (1937). On the Economic Theory of Socialism: Part Two. *The Review of Economic Studies*, *4*(2), 123-142.

Lange, O. (1967). The Computer and the Market. In C. H. Feinstein (Ed.), *Socialism, Capitalism and Economic Growth: Essays Presented to Maurice Dobb* (pp. 158-161). Cambridge University Press.

layoffs.fyi - tech and startup layoff tracker. (2026). Retrieved 03/07/2026 from https://layoffs.fyi/

Makridis, C., & Lazanski, D. (2025). Navigating the knowledge problem in the era of open AI and polycentric governance. *The Review of Austrian Economics*. https://doi.org/10.1007/s11138-025-00708-z

Marx, K. (1875). *Critique of the Gotha Programme* (English edition ed.). Progress Publishers. https://www.marxists.org/archive/marx/works/1875/gotha/

McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (1955). *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf

McKinsey & Company. (2025). *The state of AI in 2025: Agents, innovation, and transformation*. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Medeiros, N. (2022). *Big Data, Capitalism, and Socialism* [Master's Thesis, University of Porto]. FEP. https://hdl.handle.net/10216/144980

Mises, L. (1920). Economic Calculation in the Socialist Commonwealth. In F. A. Hayek (Ed.), *Collectivist Economic Planning: Critical Studies on the Possibilities of Socialism* (pp. 87-130). George Routledge & Sons, Ltd.

Morozov, E. (2019). Digital Socialism? The Calculation Debate in the Age of Big Data. *New Left Review*, *116/117*, 33-77. https://doi.org/10.64590/tt3

Motta, M., & Peitz, M. (2021). Big tech mergers. *Information Economics and Policy*, *54*, 100868. https://doi.org/https://doi.org/10.1016/j.infoecopol.2020.100868

Murel, J., & Kavlakoglu, E. (2026). What is reinforcement learning? https://www.ibm.com/think/topics/reinforcement-learning

OECD. (2026). *AI use by individuals surges across the OECD as adoption by firms continues to expand*. https://www.oecd.org/en/about/news/announcements/2026/01/ai-use-by-individuals-surges-across-the-oecd-as-adoption-by-firms-continues-to-expand.html

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge university press.

Petrova, M. (2025). A guide to the $1 trillion-worth of AI deals between OpenAI, Nvidia and others. *CNBC*. https://www.cnbc.com/2025/10/15/a-guide-to-1-trillion-worth-of-ai-deals-between-openai-nvidia.html

Phelan, S. E., & Wenzel, N. G. (2023a). Big Data, Quantum Computing, and the Economic Calculation Debate: Will Roasted Cyberpigeons Fly into the Mouths of Comrades? *Journal of Economic Behavior & Organization*, *206*, 172-181. https://doi.org/https://doi.org/10.1016/j.jebo.2022.10.018

Phelan, S. E., & Wenzel, N. G. (2023b). "Modern Computing and the Economic Calculation Debate: A Special Issue of the Journal of Economic Behavior and Organization". *Journal of Economic Behavior & Organization*, *205*, A11-A14. https://doi.org/https://doi.org/10.1016/j.jebo.2022.11.019

RIGB. (2023). What's the difference between ai and regular computing? | royal institution. https://www.rigb.org/explore-science/explore/blog/whats-difference-between-ai-and-regular-computing

Rockall, E. J., Tavares, M. M., & Pizzinelli, C. (2025). Ai adoption and inequality. *IMF Working Papers*, *2025*(068). https://doi.org/10.5089/9798229006828.001.A001

Samothrakis, S. (2021). Artificial Intelligence inspired methods for the allocation of common goods and services. *PLOS ONE*, *16*(9), e0257399. https://doi.org/10.1371/journal.pone.0257399

Samothrakis, S. (2024). Artificial intelligence and modern planned economies: a discussion on methods and institutions. *AI & Society*, *39*(6), 2961-2972. https://doi.org/10.1007/s00146-023-01826-7

Saros, D. E. (2014). *Information Technology and Socialist Construction: The End of Capital and the Transition to Socialism*. Routledge.

Stanford HAI. (2026). Economy | the 2026 AI index report | Stanford HAI. https://hai.stanford.edu/ai-index/2026-ai-index-report/economy

Stryker, C. (2025). What is agentic ai? | ibm. https://www.ibm.com/think/topics/agentic-ai

Stryker, C. (2026). What are large language models (LLMs)? https://www.ibm.com/think/topics/large-language-models

Stryker, C., & Kavlakoglu, E. (2024). What Is Artificial Intelligence (AI)? https://www.ibm.com/think/topics/artificial-intelligence

TradingView. (2025). *Biggest companies in the world by market cap*. TradingView. Retrieved 28/12/2025 from https://www.tradingview.com/markets/world-stocks/worlds-largest-companies/

Turing, A. M. (1950). I.—COMPUTING MACHINERY AND INTELLIGENCE. *Mind*, *LIX*(236), 433-460. https://doi.org/10.1093/mind/LIX.236.433

Varoufakis, Y. (2023). *Technofeudalism: What Killed Capitalism*. Bodley Head. https://books.google.com/books?id=tbqazwEACAAJ

Vedantam, S., & Penman, M. (2016, 2016/05/17/T00:01:00-04:00). NPR In *This is your brain on uber*. https://www.npr.org/2016/05/17/478266839/this-is-your-brain-on-uber