Arti cial Intelligence and Economic Calculation

Robert Gmeiner
Kennesaw State University

Mario Harper
Utah State University

1 Introduction

In a day and age when socialism is ascendant in the United States and computer technology continually improves, it is reasonable to ask if economic planning is more viable than it once was. The classic arguments against central planning, put forward by Ludwig von Mises and F.A. Hayek are now decades old and do not account for the rapid and vast changes in technology. Since that time, a small amount of research has been done on the viability of socialism with modern advances in computing power. There are arguments for and against socialism guided by arti cial intelligence (AI) and computer technology (see sections 2.1 and 3), but none of them takes a comprehensive view of the realistic capabilities and shortcomings of an AI-planned economy.

Our questions are succinct. First, we ask whether AI can plan an economy. We then ask if it can achieve results equivalent to or better than those of a free market. The rst question asks whether AI is better than humans at planning, which may be the case but this is not certain. Central planning by humans has a deplorable track record and some of its inherent failings can be corrected by AI, but AI also introduces new problems. The second question asks if AI is better than a spontaneous market order, or unplanned economy. We show that it is not. Beyond falling short of the outcomes of the free market, AI introduces new public choice problems into an economy. Some of these are magni cations of earlier public choice problems from planned economies, but others are unique. With AI, socialism is more of a public choice issue than an economic calculation issue; it is still not a viable way to organize an economy. Notwithstanding the increased prominence of public choice problems that can bring down an AI-planned economy, there is a remaining calculation problem.

economies that rely on computers tend to emphasize computing power, not artificial intelligence (Cockshott and Cottrell 1989; 1992; 1993). They focus on programs to exploit computing power, but not on systems that can produce outcomes that rival those of spontaneous orders.

Turing (1950) conceptualized a test of the thinking capability of machines. The words \thinking" and \machine" are too broad to succinctly answer the question of whether machines can think. The test to answer the question, now known as the Turing test, is that a machine can be said to \think" if it can act in a way that is indistinguishable from a human. In other words, a machine passes the Turing test if, out of sight from a human, it can fool the human into thinking that it (the machine) is a human.

The Turing test is still widely used as a benchmark in the field of AI, but it is hard to quantify and is deficient in matters of planning an economy. No single person has planned or can plan an economy, as Mises and Hayek aptly argued. For purposes of economic planning, we reframe the Turing test as follows: An AI system can plan an economy if it can produce results indistinguishable from a human spontaneous order. Reducing this to a transaction-specific test, AI can plan an economy if the incentives it produces are indistinguishable from those produced by price signals. This means that economic incentives must exist for the individuals, groups, and firms that exist in an AI planned economy and some form of entrepreneurship and innovation must take place, whether human- or machine-led.

An AI-planned economy could be thought of as \successful" if it mimics a spontaneous order. However, the counterfactual of a spontaneous order will not exist in a planned economy. Because a spontaneous order is based on voluntary exchange, one measure of success is if participants cooperate willingly and *perceive* no more coercion than in a market economy. All other measures of success or failure are subjective. If an AI system is effective enough to achieve its goals, then success is less a matter of computational success than a matter of setting goals, which is a political matter.

3.1 Coordination of Information

Rapidly obtaining information and analyzing it are major challenges for AI systems that could potentially plan an economy just as they were for central planners in earlier generations. AI has not eliminated this challenge, but has rather refined it. Data problems are now more technical and less theoretical, but they still pose major risks for an AI-planned economy.

Modern intelligence frameworks are designed to operate in belief space, or a set of possible worlds. This was in response to complications posed from applications in unstructured environments where algorithms could only obtain information from partially observable non-deterministic systems. This is especially true when algorithms respond to and make decisions involving human actors, a best judgment must be made not just on what information is presented, but on what the expected

3.1.3 Signals, Noise, and Inference

Clean and robust data are harder to obtain than complete data. Any large dataset is bound to have noise alongside signals. Major anomalies could be detected in real time, but the challenge of filtering and cleaning data still seems largely insurmountable.

Firm managers face the same problem within their firms, but their decisions have less impact than those of a central planner. Every signal obtained from data has a level of confidence based on the level of noise. Decisions made pursuant to signals in turn affect production or consumption of intermediate and final goods and the flow of more data into the AI system. A well-developed economy with many intermediate stages of production has many opportunities for errors in data gathering processes to cascade and magnify. An AI system may be able to detect if something is not right and prescribe a solution, but its success in so doing assumes that it actually knows what is going into a production process and what outputs are being produced. Without complete data, an AI system could incorrectly observe what is happening and draw erroneous inferences for other production processes.

3.2 Decision Making Algorithm

A general optimization or controls framework directs the efforts of the overall mission planning. In the case of an economic optimization problem, this mission objective is determined by the political entities who have charge over the economy.

This objective is translated into a cost, penalty, or reward function which guides the decisions towards the goals that the governing entity desires.

3.2.1 State Space

The state space representation describes the world of interest to the algorithm's objectives. This world will likely contain pertinent information on the state of consumer sentiment, consumption spending, industrial output, investment spending, firm holdings, supply chain health, and overall physical health of the populace. This is by no means an exhaustive list but serves to illustrate items that may be of importance to governing entities.

The state estimation occurs based on data received and attempts to construct the actual state of the economy. Some of this can be directly assessed or evaluated. As an example of this, industrial output data can be directly collected and investment spending can be computed given spending figures. However, some items such as physical health of a populace, general happiness, and support of a regime needs to be approximated based on the collected data.

While these terms are harder to consider, it is important to note that they play a part in a decision-maker's considerations. A well de ned state estimator for an economy must be able to predict the actual situation as best as possible in order to propagate expected changes moving forwards.

3.2.2 Control Space

The control space is de ned as the space in which the controllable inputs are considered. These controllable parameters can be policy decisions, scal policy, monetary policy, or any other type of inducement an entity can exert on the economy. It is important to explicitly consider this space and understand how action in control space at time t_i impact the state space in $t_{i+1;:::;i+n}$.

By considering the control space directly the sometimes intractable problem of solving a 2-point boundary value problem (where current state and desired states are connected by a series of control inputs) no longer must be considered. Rather, direct propagation by using search algorithms of a series of control states can directly return sets of feasible actions one can take to obtain the goal.

3.2.3 Modeling External Interactions

No economy exists in isolation. There are interactions with the world, which a national AI system cannot control. In optimizing its equations, the AI system must take into account its predictions what the world will do. Gathering information is an uncertain process in the domestic economy; it is even more so for foreign countries, even just major trading partners. This problem will not go away if all national economies are centrally planned with AI because the goals of the objective functions will vary by country.

3.2.4 Time Horizons

The goal of the AI system should determine its horizon. A focus on new inventions or research and development will be most e ective with longer horizons. Perfect information is a long-run phenomenon, but optimization is for the short run. Imperfections m326(yae9m15(will)-c 8(en)-ati6n)-3w in

3.2.5 Corrective Action

Central planning as it has existed in socialist countries during the twentieth century was characterized by inefficient use of resources and much waste. These spectacular failures obscure the fact that many resources are also wasted in markets. Around 40% of new consumer products fail in the marketplace (Castellion and Markham 2013; Crawford 1987). The failure rate for new startup companies is also quite high. According to the Bureau of Labor Statistics, between 600,000 and 700,000 new businesses are started in the United States in non-recession years. Of those founded in 2012, 79% survived their first year and 68.7% survived the second, but only 50.1% made it through five years and the survival rate continues to decline with age. This pattern of decline is stable going back to firms that started in the mid-1990s.[1] Entrepreneurs must make predictions and plan their operations accordingly and they often do not succeed. The reason markets have dramatically outperformed central planning is not because they have no waste (they do) but because waste and inefficiency are quickly corrected. The lack of profits causes unsuccessful enterprises to shut down. For any AI planning system to succeed, it must have a mechanism for detecting inefficiencies and correcting them.

To avoid these inefficiencies, both overproduction and underproduction must be penalized. Overproduction is the inefficiency discussed in the preceding paragraph, and can be prevented by a

3.3 Objectives and Outcomes

Outcomes may be prescribed by a planner, learned through algorithms, or a combination. A spontaneous order maximizes utility and pro ts of all participants through voluntary exchange. Any planned system will have some broad prescribed goals. These could include GDP growth, income equality, full employment, more new inventions, environmental quality, export promotion, production of speci c goods and services, regime stability, or any other goal. Political processes set these goals.

3.3.1 Hierarchy of Goals

In a spontaneous order, the economic system has no overarching goal. Firms have the goal of maximizing pro t. Goals exist in a hierarchy; the pro t goal is at the top, but there are related sub-goals, like recruiting top talent, investing in good equipment, producing at lower cost, and designing better products. An AI-planned system must have goals if the objective function is to be maximized by adjusting control variables.

An AI system can have a broad overarching goal and choose sub-goals to reach it, or it could have rm-level sub-goals and leave the top goal unspeci ed. The latter option will be closer to success in terms of the revised Turing test. Given that economic planning is always subject to political pressures, the former option seems more likely. There could still be a role for AI planning even if there are only rm-level goals, such as maximizing pro ts. Just as the AI planning system must have predictions for what the world will do, individual rms must make uncertain predictions about what other rms will do. Even without an overarching goal, an AI system could eliminate much of this uncertainty. Despite this potential, it seems more reasonable that a system as powerful as is needed to plan an economy would be subject to political pressures and have broad goals.

3.3.2 Planning Action and Deep Learning

Learned outcomes are an advantage of AI planning. Any action that an AI system recommends must eventually be carried out by real people. These \planners" who carry out the AI system's recommendations will act in accordance with their own goals and incentives. As they tweak the prescribed actions, deep learning algorithms can perceive what they \really want" and update recommendations accordingly. AI systems may eventually prescribe less intuitive actions to meet

3.4 Entrepreneurship and Anomaly Detection

An AI planning system must have some means of anomaly detection and correction to ensure compliance with the plan. If the control space is su ciently large and the hierarchy of goals is complete, there is no room for entrepreneurship. Entrepreneurs successfully do what nobody else has successfully done. When an attempt to do so is made, a well-designed AI planning system will detect the anomaly and, with support from government o cials, put a stop to it. If o cials choose not to act, the AI system may recognize the value of that entrepreneurial activity and cause it to expand, but will not be friendly to future entrepreneurial anomalies. Mises ([1949] 2007) opposed mixed economic systems that controlled some but not all economic activity, believing that they would inevitably degenerate into comprehensive planning. His fear did not come true in western economies, but it looks more plausible with AI planning because of the necessary completeness of the planning system.

There are three dominant strands of economic thought on entrepreneurship. Schumpeter (1943) coined the term \creative destruction" in which entrepreneurs are a disequilibrating force that introduce new innovations and displace existing businesses. Kirzner ([1973] 2013) took the opposite view that entrepreneurship is equilibrating and is only possible because existing disequilibrium leaves untaken pro t opportunities, which alert entrepreneurs seize. Entrepreneurs thus perceive a disequilibrium and pro tably correct it. Lachmann (1956, 1963) put forward another idea that entrepreneurs combine and recombine inputs looking for success. Unlike Schumpeter and Kirzner, Lachmann leaves room for research and development to be considered entrepreneurship.

the AI system, potentially cutting off follow-up work by the successful entrepreneur.

4 Economic Planning Implications of Artificial Intelligence

To say that there are public choice problems with an AI-planned economy because of the incentives of planners is an oversimplification. The nature of these problems is more nuanced. Moreover, there are issues of unintended consequences and manipulation of the flow of information that an AI system needs to plan an economy, and there is a remaining economic calculation issue.

4.1 Decision Makers and the Scope of Collective Decision Making

The scope of collective decision making expands dramatically in any planned economy. Power gets concentrated in the hands of planners. An AI system does not solve this because goals must still be set by planners. As the control space expands, the power of planners expands, magnifying the incentive problems.

Firm managers have little power because many decisions are predetermined by central planners. To the extent that they have autonomy, their choices are constrained by controlled input prices. Traditional roles such as negotiating pay rates with employees or labor unions or dealing with pressure groups that are unhappy with the firm's ethical or environmental practices become matters for planners. Labor unions, nonprofit organizations, and political pressure groups, which already officials) can rely only on programmers' explanations. They bear accountability but may lack control. Not only does the scope of collective action increase, but the locus of power shifts, along with the public choice problems.

The first public choice problem is the hierarchy of goals. Goals can be set democratically or in other ways, but they are implemented through objective functions and constraints on those functions. There is a vast technical component to goal setting. Constraints must be adequately programmed with thresholds of acceptability for certain activities and must permit decisions only in some cases. Even if the hierarchy of goals is complete, which is difficult enough, programming the hierarchy in a way that reflects the political decision is not trivial. This creates a public choice for programmers who will inevitably have discretion in writing economic planning algorithms.

The next public choice problem is collective action among AI programmers. They may or may not exist in one geographical location depending on institutional arrangements, and horizontal and vertical communications among them may be extensive or nonexistent. There will invariably be many programmers, each with different responsibilities because of the number of optimizations that need to be done and the data collection apparatuses that need continual evaluation and refinement. The outcomes of an AI-planned economy will depend on the institutional arrangements that govern programmers as they will on the individual incentives of programmers and the ability of the AI system to achieve its goals.

Another public choice problem arises with those who can influence the collecting and reporting of information to the AI system. These persons have power distinct from programmers and elected officials. Such persons could be those with control over computer network systems that oversee surveillance apparatuses and parse information for the AI system or they could be those who can successfully feed misinformation into the system. Abuses if this source of power will be difficult to detect because an AI planning system must take the information it receives as given.

4.1.2 Creating Faulty Information

Incentives exist to manipulate the AI planning system at every stage where decisions are made. If there are a large number of simultaneous optimizations subordinate to the highest, then incentive problems exist at all of these lower levels. This is an unavoidable problem because of the need to simultaneously optimize each supply chain.

Operations managers, local officials, and AI programmers for subordinate optimizations presumably have their own goals and objectives. They may attempt to manipulate the flow of information to the AI system as a way of diverting resources where they want them. These incentives can take many forms. The most obvious is if managers are rewarded for certain performance measures, like

5 Conclusion

Computers have revolutionized the economy in ways that were unimaginable even just a decade ago. Arti cial intelligence has made impressive strides. Given the power of machines to \think" and make reliable decisions, it is reasonable to reopen the economic calculation argument. This argument, advanced by Hayek and Mises, has been a mainstay of the economic refutation of socialism. Although portions of Mises's argument leave little room for an AI-planned economy, Hayek believes that can potentially be solved with su cient computing power.

The ability of AI to obtain and coordinate reliable information is still lacking to the point that planning a small economy is an insurmountable task. This is especially so when there is international trade, because AI must correctly model the rest of the world as well as the national economy. Processing data is no trivial task, even with su cient computing power, largely because noise can obscure signals as it is ampli ed through a series of cascading decisions. State-of-the-art AI algorithms o er little hope of surmounting these challenges.

The problematic issue is not speed of calculations or the inability to parse vast quantities of data; the issue is that analyses and results are not reliable enough. Human economic calculation fails in part because of inability to gather and analyze su cient information. AI-driven economic calculation fails because of the information itself. Beyond information, the ability to set goals completely enough to avoid unintended consequences is nearly impossible. Moreover, the goal-setting process is not transparent, so planners and programmers are somewhat disconnected as programmers must take planners' ideas and write them into objective functions and constraints. Finally, entrepreneurship can only exist in its fullest sense outside of state planning. Entrepreneurship is the driving force of economic growth and development, and an AI-planned economy will likely sti e it.

Socialism has always had a public choice problem because leaders were less accountable to the people and they wielded greater power than under pluralistic, free market systems. Comprehensive

Salerno, Joseph T. (1993) \Mises and Hayek Dehomogenized" *Review of Austrian Economics* 6 :113{146

| (1994) \Reply to Leland B. Yeager on `Mises and Hayek on Calculation and Knowledge'" *Review of Austrian Economics* 7(2):111{125

Schumpeter, Joseph A. (1943) *Capitalism, Socialism, and Democracy* London: George Allen and Unwin, Ltd.

Searle, John R. (1980) \Minds, Brains, and Programs" *Behavioral and Brain Sciences* 3(3):417{424

Turing, Alan J9BlenJ9Blec29s3ompuuo eS0d3c374 hinery359cI(3n)-333(n"d(3n)-3ITm [(telliDehceograms")]TJ/F74 9.96