



# Artificial intelligence and modern planned economies: a discussion on methods and institutions

Spyridon Samothrakis<sup>1</sup>

Received: 7 March 2023 / Accepted: 16 November 2023  
© The Author(s) 2024

## Abstract

Interest in computerised central economic planning (CCEP) has seen a resurgence, as there is strong demand for an alternative vision to modern free (or not so free) market liberal capitalism. Given the close links of CCEP with what we would now broadly call artificial intelligence (AI)—e.g. optimisation, game theory, function approximation, machine learning, automated reasoning—it is reasonable to draw direct analogues and perform an analysis that would help identify what commodities and institutions we should see for a CCEP programme to become successful. Following this analysis, we conclude that a CCEP economy would need to have a very different outlook from current market practices, with a focus on producing basic “interlinking” commodities (e.g. tools, processed materials, instruction videos) that consumers can use as a form of collective R &D. On an institutional level, CCEP should strive for the release of basic commodities that empower consumers by having as many alternative uses as possible, but also making sure that a baseline of basic necessities is widely available.

**Keywords** Artificial intelligence · Planning · Reinforcement learning · Adaptive systems

## 1 Introduction

Among the many efforts to envision a system that will help us move beyond our current state of affairs, arguably one of the most prominent cases involves the use of computers to help guide production, which we will term computerised economic planning (CEP) from now on (e.g. Cockshott and Cottrell 1993). While the exact system in place varies from author to author, the main essence is the use of computers to calculate what commodities to produce, doing away with (or partially substituting) the market. The debate on the (conceptual and theoretical) feasibility of such reforms goes back to the early 20th century, with market proponents like Hayek making the crucial link between knowledge, computation, and economic viability (for an overview, see Yeager 2004). They envisioned the market as an agent, acquiring knowledge, performing calculations, and acting in the most rational way possible. Throughout the years, debates and proposals have been inspired by and/or have touched upon other numerate disciplines like cybernetics, game theory,

optimisation, complex systems, machine learning, and statistics.

Arguably, the applied cutting edge of the fields that have partially contributed to the debate is currently being studied under the broad umbrella of artificial intelligence (AI). Contrary to what popular media or corporate marketing propagate, AI is nowhere close to replacing human labour, as quite a few problems that originally seemed easy elude today’s state of the art. The archetypal failure case is self-driving cars; they have been promised for years, but it looks increasingly unlikely that fully autonomous driving can be solved without resolving the problem of (almost) replicating human intelligence. The debate as to why AI proved harder than once anticipated has been going on for 50 years and is still raging, but there seems to be broad agreement across two themes. The first one is that most knowledge in the world is tacit, i.e. it does not exist in a form that can be trivially “algorithmised” and mechanised, as it cannot be verbally or mathematically expressed. Even worse, the hopes that we had that certain fields of AI (e.g. machine learning) would be able to capture that knowledge have so far been proved unsuccessful. The second theme is flat-tailed distributions (Taleb 2007); there is a massive amount of marginal events that can happen when trying to predict even the most trivial of future outcomes, events that are hard to

✉ Spyridon Samothrakis  
ssamot@essex.ac.uk

<sup>1</sup> Institute for Analytics and Data Science, University of Essex, Wivenhoe Park, Colchester, Essex CO4 3SQ, UK

conceptualise early on (e.g. bee infestation when teaching a class). These “unknown unknowns” are far too many to list, making it difficult to find algorithmic solutions in open systems. The scale of these issues has made some researchers sceptical of AI promises, and publications are emerging that try to tone down the overhyped rhetoric (e.g. Narayanan 2019).

Although fully automated pipelines are not yet available and machines that can generally replace humans may not be available for a long time (Allyn-Feuer and Sanders 2023), AI is widely used in the business world as a form of advanced analytics. It is clear that most private businesses attempt to predict the future; energy firms calculate future energy needs, factories fix orders years ahead, and massive conglomerates have the ability to source and sell billions of different products. In this sense, the multiple institutions of capitalism, from state departments to highly exploitative megacorporations, have proven both resilient and adept in some form of prediction and evaluation, which some would call planning (Phillips and Rozworski 2019). Internally, these “planning” systems employed look nothing like the planning systems envisioned by current socialist thinking (e.g. see Cockshott and Cottrell 1993). Capitalist institutions have created a very convoluted symbiosis of human and machine systems that together allow them to peer into the future as clearly as one can, while their immense political power often makes them immune to glaring mistakes and abuses. Given widespread proliferation and hype around AI methods, it is somewhat surprising that CEP proponents have not been involved heavily in technical developments in the field of AI. One of the major issues modern work on CEP suffers from is that it tries to link itself to outdated (but excellent) ideas from the 1960s, mostly around celebrated figures like Kantorovich and Beers. Though it is not clear why, the suspicion here is that this was the last time that cybernetics and aligned fields were concerned with anything different than serving moneyed interests. Within AI communities, money and markets have a quasi-mythical status, so any discussion on CEP might be seen as some form of superstition, which in turn makes the literature produced by these communities hostile to CEP ambitions from the outset. What we are working towards in this paper (and using Boyd (1987) and Beer (1979) as inspiration) is unearthing what we know about agents from AI and try to apply it on designing new organisational systems, in the same spirit as they did with thermodynamics and early cybernetics. In this sense, the closest analogue one could give between CEP and AI is that relationship held between AI and robotics. Robotics is not an application of AI methods and is a field of its own with its own conferences, research venues, and grant structures. However, the cross-pollination between the two fields is significant, with results moving between fields continuously.

The line of thinking in this paper and the proposal put forward encode some basic assumptions about the prospects of both technology and humanity and our capacity to learn from nature and leans on the pessimistic side; one can only hope for small incremental steps and quite a bit of back and forth. We see our version of CEP as a transitory institution, whose need would arise from societal pressures that are beyond the scope of this article, not as a model for a post-capitalist economy, as it is made to co-exist with current structures, parallel to the market but would allow for the flourishing for a much wider group of the population than what we have now. Such an institution would be of an “equalising” bend and would provide alternatives that some of the elites would find palatable (and thus be more inclined not to end the world in direct opposition). A good analogue would be the early Christian church. Christianity offered a moral code and an institution that helped curtail excessive elite male power. As an example, the extreme Roman imbalance in the power of sexes is gradually transformed through monogamy, the banning of abortions (as a counterweight to female infanticide), and generally a set of less liberal but protective measures (Stark 1996). It took almost 2000 years before humanity was able to set up procedures in place and before technology would catch up to the point where we could rethink these practices. As a group, the clergy and its beneficiaries (the core members of this new institution) seem to have had a ratio of 1 in 10 (Grant 1970) with the general population of the church. We envision a CEP institution of an equally transformative structure, one that would change the way we produce and consume, with a focus on constraining the unrestrained freedom of the markets that control production and distribution, while also juggling between competing interests and adversaries, without becoming overbearing and overly controlling.

The rest of the paper is organised as follows. In “[Planning, reinforcement learning, and AI](#)”, we briefly introduce some central notions in modern AI. In “[Labour theory of value, markets and reinforcement learning](#)” we further develop the AI framework, this time to portray a simplified model of our current economies that partially links them to AI. We discuss planning, how it was carried out in the Soviet Union and its links to AI in “[Problems in traditional CEP](#)”. Following the “lessons learnt”, we hypothesise how a planning institution would look like in “[A planning institution](#)”. We conclude with future research directions in “[Final thoughts](#)”.

## 2 Planning, reinforcement learning, and AI

Reinforcement learning formalises an agent-centric view of the world. In its most basic setup, RL deals with a tuple known as the Markov decision process (MDP). 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  $a \in \mathcal{A}$ . A transition function  $P_{s,s'}^a$  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  $R_s^a$ , which denotes the amount of “pleasure” an agent receives by performing this action in the specified state. Finally, a discount factor  $\gamma$  denotes how the agent treats rewards time-wise; in its extrema, if it is set to 0 the agent only cares about immediate rewards, while if it is set to 1 the agent treats all rewards, no matter how far into the future, the same. In the classic RL setting (think a simple video game or maze), the agent traverses the world using initially some quasi-random policy  $\pi(s, a)$ . The goal of an agent is to discover a policy  $\pi^*(s, a)$  that maximises the average sum of long-term rewards,  $V$  as in Eq. 1:

$$\pi^* = \operatorname{argmax}_{\pi \in \Pi} V_\pi. \tag{1}$$

Each MDP state has a corresponding value  $V$ , which denotes this average sum of long-term rewards. More precisely, and following Sutton and Barto (2018), one can define  $V$ -values as in Eq. (2)

$$\begin{aligned} V_\pi(s) &= \sum_{a \in \mathcal{A}} \pi(s, a) \left[ R_s^a + \gamma \sum_{s' \in S} P_{s,s'}^a V_\pi(s') \right], \\ V_\pi(s) &= \sum_{a \in \mathcal{A}} \pi(s, a) R_s^a + \gamma \sum_{a \in \mathcal{A}} \pi(s, a) \sum_{s' \in S} P_{s,s'}^a V_\pi(s'). \end{aligned} \tag{2}$$

If we assume that rewards are given only when landing or leaving a state (i.e. they are not conditioned on actions, a further simplification), we get Eq. (3):

$$\begin{aligned} V_\pi(s) &= \sum_{a \in \mathcal{A}} \pi(s, a) R_s + \gamma \sum_{a \in \mathcal{A}} \pi(s, a) \sum_{s' \in S} P_{s,s'}^a V_\pi(s'), \\ V_\pi(s) &= R_s + \gamma \sum_{a \in \mathcal{A}} \pi(s, a) \sum_{s' \in S} P_{s,s'}^a V_\pi(s'). \end{aligned} \tag{3}$$

We can combine a predefined policy with the transition matrix to get Eq. (4):

$$V_\pi(s) = R_s + P_{s,s'}^\pi V(s')_\pi. \tag{4}$$

*Planning*, in this setting, refers to an agent that has somehow internalised the MDP (i.e. it knows everything there is about the world), spends some time thinking really hard and executes in the real world whatever policy it came up with. In large state spaces, these often take the form of model predictive control (or rolling-horizon control), where planning and replanning take place as often as computational resources allow. The prime modern example of such a method is Monte Carlo tree search (MCTS). Reinforcement learning (RL) refers to an agent that starts without any a priori conception of the world and traverses it until it can

formulate a policy. Note that RL is an overloaded concept, as it can refer to the “RL problem” (as in this section) or methods from the RL literature that attack the RL problem (e.g. Q-learning).

### 3 Labour theory of value, markets and reinforcement learning

Following the RL abstractions above, we can give an RL take to certain portions of Marxist thinking. Let us define a set of commodity concepts as  $\kappa \in K$  and denote each concept as  $\kappa_0, \kappa_1, \dots, \kappa_n$  as a set of actions which an agent can take to bring a commodity to life. In other words, an action is the whole procedure of coming up with a concept and helping make it real; if we would have used a verb here (to emphasise that actions are things the agent does), it would have been “conceptualise and materialise” or “conjure”. A manifestation of the commodity concept in the world can be described as a commodity  $c \in C$  and we denote each commodity as  $c_0, c_1, \dots, c_n$ . If we define rewards in monetary terms (e.g. 10 dollars for every book sold), our  $V$ -values become almost synonymous to exchange values. Other definitions of reward will give us other values (e.g. use values), but the abstraction in its core remains unchanged. The agents “conjure” ideas and throw them to the market, which then provides a reward, the average sum of rewards which we call the value.

The discussion above makes it easy to understand various propositions put forward by the labour theory of value, like, for example, that all value is derived from labour and nature (Marx 1875)—the advantage of this simple reformulation of the problem is that it allows us to add a series of recent advances in RL, with agency given to the abstract capitalist. Notice the absence of anything apart from commodity concepts and commodities in our formulation; labour is simply another commodity, as is nature or anything else, to be used (with the help of other concepts and commodities) to maximise value (or, equivalently, the sum of long-term rewards). Given this, we can analyse one of the most basic debates on labour.

#### 3.1 The agent traversing this MDP is worth everything

The agent deserves all reward. Any reasonable subjectivist would say that yes, labour and resources are in the mix, but even if the agent (i.e. the capitalist) pays labour very little and pockets all profit, that is fine, they came up with the “commodity concept” and executed the policy. It is the moral thing to do (albeit of maybe of a Calvinist bent). Labour in this formulation does not even have agency; it resembles a

pool of extremely reconfigurable robots. It deserves nothing outside the value-capture process (i.e. the agent's policy).

### 3.2 The agent is stealing everything, and labour and nature are the source of value

A socialist would tell you [paraphrasing Marx's Critique of the Gotha Programme (Marx 1875), and going back all the way to the diggers (Winstanley 1649)] that all wealth comes from labour and nature; the fact that an agent found itself in a position to come up with "commodity concepts" does not mean that they should be earning anything more than what they would have been paid from their own socially necessary labour time. In this view, the concept and the policy are worth nothing. They are obvious social constructs; the agent somehow inherited them; the only reason the agent is an agent in the first place is because of primitive accumulation and/or blind luck. The difference between the value of the agent and the value that goes to labour in the form of wages is theft, and it is straightforward to calculate how much value the agent is stealing from labour.

Following the collapse of the Soviet Union, (a) is the default worldview. We see economy as concepts and policy (i.e. what to make and how to make it), and all value derives from them. It is really hard to take position (b) seriously, as a naive causal reading (which most humans do by default) is that labour and nature cause value, absent of a concept and a policy, and this is evidently not true. Labour alone without a concept is obviously just a waste of time and effort. What position (b) really says is that more or less anyone can be an agent, the problem is not that hard, it is just that the vast majority of people will never have the opportunity. The counterfactual is that if you sample a single child of any reasonable ability, you can turn them into CEO pretty quickly.

## 4 Problems in traditional CEP

The historical roots of planning can be traced back to early socialist thought and are largely inspired by war economists like Neurath (Cartwright et al. 1996). They follow more or less the paradigm of the previous section. The very fact that planning was attempted before the full development of computers is an impressive feat on its own, so it is worth a brief remark. The Soviets attempted to plan only part of their economy, through a single institution, GOSPLAN, using what they called "planning using material balances". GOSPLAN planners would keep statistics, party directives, and goals and try to match production and consumption with intermediate products. This procedure was carried out by hand, was tedious and took months of back and forth to complete (Montias 1959). What the planners were trying to do (by hand) seems very similar to solving a very large

input–output problem, but with some inputs fixed to correspond to political demands. More specifically, the overall setup can be described as:

$$x = Ax + d. \quad (5)$$

In Eq. (5),  $d$  is the final demand for each product, projected at some point in the future.  $A$  is called the technical coefficient matrix, which captures what is needed to make a product (e.g. to make a cake one needs 0.5 kg of sugar, 2 kg of flour, 3 eggs, 200 g of butter, an oven, and 2 h of labour). The model is rather simple, but, in general, it does serve as a good starting point for planning discussions, and it encapsulates the ideas discussed by a large number of socialist authors (Cockshott and Cottrell 1993; Sraffa 1960; Leontief 1986; Castoriadis 1974). In its modern incarnation, this form of planning is termed "input–output" table analysis. In general, it suffers from all the problems that plague AI, possibly exacerbated.

### 4.1 Input–output table analysis is not planning in the formal sense

Let us revisit Eq. (4). In matrix notation, it can be rewritten (see Barto and Duff 1993) as in Eq. (6):

$$V_{\pi} = R_c + P_{s,s'}^{\pi} V_{\pi}. \quad (6)$$

This is highly analogous to a traditional input–output system; we only need to change the notation to  $x = Ax + d$ , with  $A$  corresponding to the combined policy and transition matrix  $P_{s,s'}^{\pi}$ ,  $d$  to the rewards  $R_c$ . Note that in the formal sense, we are not doing planning. If one contemplates an extremely large but finite space of concepts, planning would effectively be discovering new concepts. All we do is solve the input–output system for a specific set of predefined concepts that were discovered as valuable through the market. It is also not a computational problem; modern solvers can solve approximately for trillions of states (Anthony et al. 2017), provided there are regularities, which are bound to be present (e.g. different shoe sizes point to latent variables of certain age groups). Planning would mean discovering whether or not a concept  $k_i$  produces any value and how to build new commodities. If we decide to use input–output tables given the quantities of an economy in 2023 (when this paper is written), all we are going to have in 2030, after quite a bit of effort, is a perfect 2023 economy. At best, all an optimal economy that is (naively) based on input–output tables can do is play catch-up with a neighbouring capitalist economy—it is unable to conjure new valuable concepts! As the capitalist economy continuously creates new concepts and new production methods, the socialist economy would try to create an optimal version of the capitalist past. There have been multiple counter-arguments (Cockshott and

Cottrell 1993; Adaman and Devine 2001; Albert 2004; Hahnel 2005) to the idea that planning cannot support innovation. To some extent, all authors separate innovation from economic efficiency. Cockshott and Cottrell (1993) positions innovation outside the planning system and discusses it more in cultural terms, as a rate of adoption of new scientific discoveries in the general production process. Adaman and Devine (2001) take a view that innovation should come from committees and councils, as a normal part of operations. Parecon Albert (2004) takes a similar view. Hahnel (2005) claims that innovation can take place at the council level (in our model this would make the council the agent): quoting: “There are strong incentives for worker councils to search for innovations that increase the social benefits of their outputs, or reduce the social costs of their inputs since this would increase the worker council’s social benefit to social cost ratio. Raising the social benefit-to- social cost ratio makes it easier for the council to get its proposals accepted in the participatory-planning process, can allow workers to reduce their efforts, can permit them to improve the quality of their work lives, or can raise the average effort rating the council can award its members”. The problem with such proposals is that they are often afterthoughts and responses to outside criticism. Whether a council has incentives to innovate without the appropriate support and the relevant politics, in the very best of cases, debatable.

If CEP planning is not planning (in the RL/control engineering sense), then what is it? The best way to describe intuitively is by bringing forth the equilibrium state of the market to the present, rather than letting the market discover it (and possibly correct it). If the rewards are not monetary, but rather correspond to labour hours and commodities (something termed “accounting in natura”) things become somewhat more complicated, as there is not one unit of accounting for rewards, but the basic premise remains. Although we are not aware of a formal term for this, we would prefer to call it “value mixing”, and this is the term we will use from now on.

## 4.2 Control is not measurement

There is a tendency to confuse the concept of a product with the product itself, i.e. mixing up the map with territory. Given that if you plan the way we describe, a planner’s actions are limited to requesting the development of certain concepts, there is no guarantee that these concepts will manifest as they are truly desired. This is an instance of Goodhart’s law (Teney et al. 2020; Goodhart 2015). A loaf of bread satisfies a specific human need. If you ask a factory to create 100 loafs of bread, there are enormous incentives to cut corners. Since no exchange is recorded, there is no reason for the quality of the product to plummet. As a response, a CEP institution would try to provide guidelines that are

exceptionally detailed, resulting in even more sophisticated forms of cheating and an overbearing bureaucracy. Similar problems occur under capitalism when monopolies are dominant, but with CEP in place, there is the risk of them becoming endemic. There are also dangers of premature “overoptimisation” over unknown quantities, although one would presume that common sense from planning institutions would prevent this from taking place.

## 4.3 Nonlinearity of the input/output mapping

In general, simple linear systems might not be able to capture the properties of production; for example, if a factory needs 225 g of sugar and 225 g of butter and 225 g of flour and 2 h of labour to make one cake, it does not imply a linear relationship between flour/sugar/butter/labour and the number of cakes that can be made. There might also be a physical upper limit, and you might also run into problems with lower volumes, as the machine could potentially only be operated under certain conditions. In effect, each producer should not communicate input–output coefficients, but functions that describe how each factory (or more generally, production unit) operates. This problem is addressed in Samothrakis (2021), but its complexity is not trivial and its solution increases computational demands. This makes it hard to merely shadow the existing economy and come up with control rules—one has to actively experiment on the edges.

## 4.4 Corporations probably do a mixture of planning an active RL

Corporations themselves do not (most of the time) plan in their everyday operations—projecting value is not planning, as discussed above. They follow an iterative procedure of designing a product and then releasing it to the market. The planning stage has to do with the concept “discovery”, i.e. coming up with new methods for steel production. Planning, in this sense, occurs only during the initial R & D that needs to be done to bring a product to market. Any further refinements are closer to RL, as there is a constant feedback loop between the market and the corporation. The “obtain user feedback as quickly as possible” methodology is widely adopted by modern venture capitalists and business gurus (Thiel and Masters 2014). The constant mining of user data points to a similar direction, but it is hard to predict where the value comes from. There is a counter-argument to be made that capitalism is unable to conjure the future at this point, and we are all stuck in the perpetual 90s, but one would presume that the goal of CEP would be to move us away from this.



#### 4.5 The modelling itself is extremely unrealistic

The modelling we have done so far is a crude abstraction. Using it to do any kind of planning would cause problems similar to the ones described in Scott (2008)—the action abstraction (“commodity concept”) and the state abstractions (“commodity”) in no way reflect an actual process of designing and producing commodities, which is extremely dynamic and unpredictable in nature (Kornai 1979). Even as an abstraction, the modelling fails to take into account time, failures in measurement, antagonisms, and a host of other things that would make it hard to achieve anything useful. Furthermore, production processes tend to have multiple outcomes, which in general would require multiple simulations running concurrently to help model them. It is also not clear what one would use as proxy for demand. Under the current system, money conveys all information about a commodity, and hence the translation is easy, but one can easily imagine a situation where multiple different objectives need to be met. One pair of jeans can be converted into 100 eggs—which one is more valuable? The amount of research that is needed to build a CEP system should not be underestimated. It is likely that a whole new branch of science, possibly as a subfield of economics and computer science, would need to emerge to tackle these problems and understand their general patterns.

#### 4.6 Function approximation, catastrophic forgetting and plasticity–stability

On a technical level, systems become difficult to solve when faced with a set of problems known as the “deadly triad” (Sutton and Barto 2018): learning from observing others (off-policy learning), learning approximate policies (function approximation) and bootstrapping (learning from distal reward signals). None of these are showstoppers for humans, but tend to pose significant problems for machines. Arguably, the most important of these problems is function approximation. Roughly, because the number of all the potential concepts one can come up with is very high (in a game like chess this would correspond to the number of actions times the number of states), one needs to find an approximate solution. As one keeps discovering new concepts, measuring their quality, and concentrating on them to produce value, old concepts are forgotten. Game-playing RL agents use a number of methods to go around these problems (mostly around re-play buffers), but overall these are workarounds born out of necessity, rather than principled solutions. The problem of what to remember and what to forget can be thought of as another instance of the plasticity–stability problem (French 1999).

#### 4.7 Demand is created as much as discovered

The common assumption (in modern RL, but quite often in economics as well) is that demand for a commodity is a by-product of human needs, for example, you are (naturally) bored, so you are looking to play a game. This attitude mimics late Soviet efforts to introduce computers (and is not captured by our MDP formulation). As Adam Curtis puts it (Curtis and PoliticsJOE 2022) “...thought okay we’ll have computers we’ll get computers in and computers will rescue communism because it’s the modern way but I found this interview with this woman who worked for a thing called GOSPLAN which was the planning organisation that planned everything for everyone that’s in the state and yeah they had computers and the computers predicted to them what people would want and so the people the computers predicted that people wanted stacked heels uh this was at the end of the 70s I mean the end of the 80s which I think is probably a bit late given that it was fashionable in the 70s I think in Britain but anyway they thought it was going to rescue them because that’s what the computer said but actually by the time they’d got around telling the factories how many to produce and how to build a stacked heel it had gone out of fashion so whatever they tried didn’t work...”. Advertising, implicit trends, and word of mouth create demand for something; reward (or, interchangeably, demand) is often the product of the agents’ mind as much as of the world around it, and is probably imbued with strong self-referential idealistic properties.

#### 4.8 Alienation and feedback mechanisms

There is an inherent tension between production and consumption; while consumption is seen mainly as a pleasant activity, production is much less enjoyed (or, at best, its enjoyment has ascetic qualities). Marx famously has drawn a more intricate distinction (see here for a review Klagge 1986) between the “realm of necessity” and the “realm of freedom”. Whereas the realm of freedom is harder to define, the realm of necessity would include all these activities that one gains no pleasure of any kind from, but they have to be done to maintain individual, family, and societal structures. The more typical response historically is one where we would recognise the existence of a realm of necessity as a necessary evil and would like to transform our lives by minimising human exposure to it, i.e. cutting down working hours (thinkers in this camp would include an extremely heterodox group of politicians/economists like Nixon, Stalin and Keynes). A second group of thinkers (Bookchin 1982; Graeber and Wengrow 2021) would like to transform work into play, that is, merge these two realms. Under CEP one can trivially imagine scenarios where product quality is monitored, but the whole CEP has an obvious blindspot on

the conditions of labour, i.e. what goes on inside a factory. The problem is hard to solve, and one could argue that modes of production are mostly ways to externalise the realm of necessity to third parties. Solutions to are again not trivial; when the Soviets instituted reforms that would make quality control a part of everyday processes, the reforms disrupted production batches to the point of almost collapse (Zubok 2021). In the less top-down and more worker-led conditions of Yugoslavia, similar patterns were observed (Miljković 2017), with quality lagging behind the more coercive analogues in the eastern and western blocks.

## 5 Goals and principles

Given the above criticisms, what concepts could one directly draw from current AI developments to help with CEP? If any attempt to model processes accurately is bound to fail and any abstraction hard to control (points 4.2, 4.3, 4.4, 4.5 above), we would like to do as little (but extremely thorough) modelling as possible and allow for economic agents to decide on their own, i.e. give them better means to build and decide on the ground. CEP should avoid having to deal with details and focus on “the big picture”, where the policy is easy(ier). We would also like product quality to remain high, so the feedback mechanism should come from the users of a product (i.e. not through workers’ councils or voting mechanisms). Functions like quality control should be external to the production system, but at the same time necessary labour should be cut to a minimum, and possibly be abolished when and if technology allows for it, allowing only fun activities to come in (in response to 4.1, 4.7 and 4.8). Finally, we also would like to address the problem of achieving widespread innovation (4.1). We address these three groups of problems in detail in what follows.

### 5.1 Commodity power

The MDP setup presented above assumes that the reward structure never changes— the same things that made one happy yesterday would make them happy tomorrow; transferring this to commodities would entail a static preference function for the whole population that is easily available. In traditional RL, a concept that has recently been introduced that is capable of helping agents adapt in different reward regimes is one of (Salge et al. 2014) or power (Turner et al. 2021), defined in Eq. (7):

$$\text{POWER}(s, \gamma) = \frac{1 - \gamma}{\gamma} \mathbb{E}_{R \sim \mathcal{D}_{\text{bound}}} [V_R^*(s, \gamma) - R(s)]. \quad (7)$$

Given a certain discount factor  $\gamma$ , the power of the value function in a specific state is defined as the expectation of

the optimal value function minus the reward in the state, under all distributions  $\mathcal{D}_{\text{bound}}$  of possible MDPs. In other words, a state is powerful if it is useful under different reward regimes. The prime example of this is “not dying”— it is a prerequisite for any form of happiness. Another way to see power is through the lens of premature optimisation avoidance. Agents prefer to be in conditions that would allow them to maximise their future freedom and choices, as they are unaware of what these choices would have to be.

If we apply this principle to our economy MDP, we end up with the simple notion that economies should create empowering commodities. In other words, good commodities (which in our case correspond to states) are those that can be used to create as many other commodities as possible, similar to basic commodities introduced by Piero Sraffa (Roncaglia 2009). An example of a powerful and rewarding commodity would be water, as it is necessary for life and can be used no matter what we want to make. A pre-packaged meal with duration of 1 day is a disempowering commodity, as it can only be used to be consumed (i.e. it has one use). In this sense, capital is raw power, as it can be transformed into any other commodity, something that partially echoes recent lines of thought (Nitzan and Bichler 2009). Empowering commodities, combined with labour power, can form the basis for an extremely versatile system of production. The end result would be an economy with shorter supply chains and more local production. Home or communal mills would be preferable to direct shipping of bread from large bakeries, raw coffee beans, and coffee roasters preferable to ground coffee, and so on. In the same vein, the home/communal devices that help with manufacturing are better off being designed with interchangeability and compatibility in mind (e.g. same battery sizes and as many similar components as possible) to allow for home innovation and assembly. This more localised mode of production would be in line partially with past work in designing economies on a smaller scale (e.g. Schumacher 2011), with our focus being more empowering individuals to take more control over their lives.

### 5.2 Decentralisation of production

Looking at the market MDP, in a market economy, the only people that can traverse it (i.e. search for value) are capitalists, a small subset of humanity that has already accumulated considerable rewards. In a naive CEP solution, the MDP could be traversed by a central planning committee or an algorithm. However, again this gives agency to a limited number of individuals. Ideally, we would like as many agents as possible to create as much value as possible, with the CEP helping with overall coordination. That is, we would like the decisions about what to be created and how this creation is to happen to be as distributed as possible. However, notions

of rationality on an individual level can be catastrophic on a global level—this is known in game theory as the “price of anarchy” (Koutsoupias and Papadimitriou 2009). If one agent decides to hoard a rare and extremely valuable commodity, they might become rich—if every other agent tries to do the same, they will all starve.

To solve this problem, we would have to step up and complexify our models a bit. In RL, different models capture different formulations of a problem. In the simplest case, a deterministic MDP is one where there is a single agent (or multiple agents in perfect coordination, all sharing all information) and they make a decision on what to make, with the environment being deterministic. A step up in complexity is an MDP, where the environment itself is stochastic, this translates to “softer” constraints on how commodities can be made. A partially observable MDP (POMDP) (Kaelbling et al. 1998) is one in which the commodities produced cannot be observed directly, but the agent that traverses it needs to work out how to form beliefs over every potential commodity. This is much closer to reality, since capturing all the statistics that describe a commodity exactly (in terms of specifications) is virtually impossible. If one starts adding other agents to the mix (or alternatively but equivalently, the agents that take part in the decision-making process are not in perfect sensory, affectory, and reward harmony, (e.g. think of a beehive), one moves to DEC-MDPs. In the DEC-POMDP (Oliehoek 2012) setting, agents need to form beliefs on the nature of a commodity (akin to a POMDP), but they are also found in widely different parts of the state space, all with the same goal. In Markov games (Littman 1994) and extensive form games (Roth and Erev 1995), multiple agents potentially compete for resources. Assuming that we do not want to pay the price of anarchy, but find a way to create value and share as much information as possible on how to do it, one would need to create the regulatory environment that would push towards DEC-MDPs or somehow approximate it. This, in practice, would mean creating coordination tools, having public board, and common educational programmes on how to build things, while at the same time advancing a prosociality culture. This is not the first time this observation has been made, and though we have arrived at it through different means, it has been proposed quite frequently as part of a wider package of methods that would help with democratisation (e.g. Castoriadis 1974) of production.

### 5.3 Equivalence of process

Creating the same final commodity requires a production process that is extremely well timed, confined, and generally roboticised. Although some powerful intermediate products would need to conform to this, so as to make interchangeable components, for final products, one might have

to follow a more “Goethean science” or phenomenological approach (Bortoft 2012). We will try to “harvest” commodities, rather than manufacture them, with each commodity ending up slightly different. One can think of cookbooks, instruction manuals, education, and other means of that would help in maximising the autonomy of an agent—no two cakes would be the same. This creates a stark dichotomy on commodities and production processes that are designed in a way to make the final product as uniform as possible (e.g. tools), and ones that are “grown” domestically or grown in the fields (e.g. potatoes). 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).

## 6 A planning institution

So far, we have described three key principles, commodity power, decentralisation of production, and equivalence of process, as the central themes of a modern CEP process, which will be built on top of value mixing. We still have issues related to plasticity and stability, which we have not addressed and which would probably need to be resolved through other means, as they have also not been adequately attacked in the AI literature. Topics around these themes (e.g. catastrophic forgetting, incremental learning) remain largely unsolved, and it is not clear when we are to expect progress in these areas. What we would like to achieve is create an institution (in the abstract sense) that would turn these principles and goals into reality; it should help design and discover empowering commodities and release them, collect educational examples on how to use them, and teach the general population on a process of improving these commodities and reporting back on their everyday use. It sits orthogonally with the principles we discussed above, complementing them by attacking the stability/plasticity of the whole system as it unfolds in time. We can identify the following main strands of work such an institution would have to do.

- (i) Value identification: identify where value is created and capture it, in support of commodity power, decentralisation, and plasticity.
- (ii) Value protection: prevent value from being destroyed, in support of commodity power, decentralisation, and stability.



- (iii) Value mixing: help bring future value to the present, in support of reward maximisation, decentralisation, and equivalence of process.
- (iv) Value ethics: distribute commodities fairly, in support of decentralisation, and also ensure that value is not extracted in ways that would be immoral.
- (v) Meta-value: evaluation of the programme.

The five above strands are in constant dialectical battle with each other. Supporting value creation (i) might require stopping old products (as discussed in the previous section, akin to the plasticity–stability dilemma in artificial learning systems); hence, it conflicts with (ii), stopping value from being destroyed. Strand (iii) is more straightforward and is what most socialists recognise as “planning”. Once we have decided what we will do, we tell each industry how much to create. Strand (iv) must be well informed of the methods employed by strand (iii) and stop anything deemed unethical, irrespective of demand. Strand (v) creates coherence in the entire CEP institution by helping it self-evaluate. Central to any CEP institution is the collection of data; the institution would need to collect a very wide array of data around production and consumption, which would range from consumption habits to what components of devices break down and how often, post-process them, and release them back to the public, and the way these strands deal with data would be crucial in their effort to achieve their mission.

One can almost immediately observe that whoever controls the institution of CEP could exercise extremely strong influence on what and how it is produced. Ideally, any institution that we would like to make accountable and democratic should operate through a jury-like system inspired by Aristotle’s *Politics*. We expect nobody to be permanently employed by the implementing organisations of the CEP institution, with secondments from other areas lasting 5–7 years, while also having a large number of citizens drawn into the institution by lot. The general idea is to avoid someone using CEP as a springboard to elite power, but also to give ample opportunities for participation. We also expect the larger bureaucratic ecosystem of the state to support CEP and feed it with ideas and people.

### 6.1 Value identification

In analogy to neural processes, this strand represents the plasticity component of the overall institution of CEP. It searches the world for valuable, empowering, and easy-to-create products. It champions new ideas that are ready for production and pushes them to strand (iii) for further inclusion within production processes. The data it holds are of a research bent, and the modelling from this strand would be on identifying unfulfilled user needs. Home visits can help to understand how old commodities are used, what new

commodities to introduce, setting procedures, competitions, and artistic representations of the future. It would need to be tightly integrated with corporate R & D and university laboratories—similar thoughts were put together by Nieto and Mateo (2020). This strand would resemble a cross between a venture capitalist, an entrepreneur, a deep tech funding body, and a pure research outfit, but with different goals than profit extraction.

### 6.2 Value protection

This strand would identify which commodities are responsible for the minimum quality of life and act as the “stability” part of CEP. It would oversee multiple producers and develop commodities ready for direct consumption, but also basic tools. These would include items such as shoes, basic hygiene items, and basic food provision. Items of this type would have to be delivered directly to consumers, ideally through mechanised means, at regular intervals. Strand (ii) would not just identify what is currently needed in terms of commodities, but what might be needed in the case of emergencies and other unwarranted events. Constraints in production would likely come from this part of the organisation. The value protection strand should have a mission that aims to stop boring parts of the economy from disappearing and/or overinvesting in exciting new products. It is the conservative part of CEP, and would be in constant conflict with the “value identification” part. It would veto new products by eating resources for older products and promote tradition and stability over new and untested products; a place where Chesterton’s fence (Sutherland 2016) would rule supreme.

### 6.3 Value mixing

Once what is to be made is identified, it would need to be produced. Strand (iii) of the institution would act on this, do the calculations, and collect user feedback. In RL terms, once the actions have been identified, by the value protection and value creation aspect of the institution, the value mixing part would operationalise these decisions. The constraints as to what needs producing no matter what the current consumption targets can be calculated through collective statistics from the general population, through future projections. Properly doing this is not trivial, as production and consumption is a concurrent process, but the basic mathematical tools are there, and are to be drawn largely from optimisation and RL (e.g. Badia et al. 2020), where Monte Carlo under function approximation solutions are the norm. The value mixing strand would handle the core calculation operations and would be what the CEP community currently recognises as planning. It would also need to be the place where large databases of everything that is known about production are kept and regularly updated, while also being made widely

available through easily accessible websites and announcement boards. Overall, this strand is the most “technical” strand of the CEP institution, and is more concerned with the “hows” rather than the “whats” of production. It is a cross between a library, an accounting strand, a corporate planning department, and technical journalism.

#### 6.4 Value ethics

Fairness for a CEP institution is not trivial. The needs of different people in the population will vary widely depending on highly contextualised situations. Differences in ability, talents, background, and personal preferences are difficult to incorporate within a framework that places too much emphasis on the average case (Lee et al. 2021). The easiest way for CEP to be fair would be to create conditions in which there is no shortage of any type for anyone at all times. Every commodity is available in abundance, and any issues that might arise are more a distribution problem. However, this situation is very unlikely. Ideally, given that most commodities will have an exceedingly high degree of power, problems of fairness would be minimised. For example, provided that there was a tool to cut potatoes in various shapes, there would never be an issue were a group (“people who like their chips cut thin”) could complain that the planning system is unfair in terms of only providing another group of consumers with access to products (“people who like thick chips”). This strand will also have to work on problems that arise from production processes (such as certain forms of dangerous labour, unethical, and/or extremely extractive use of animals). Mathematically, solutions to such problems have been widely studied (e.g. additional constraints are added to the reward function, see Garcia and Fernández 2015 for a review), but issues remain in compiling all these constraints in a case-by-case basis and having a deep enough understanding of the production processes and the tools used. Another major issue that arises when discussing fairness is the problem of allocating time. The MDP we describe is an abstraction that does not take into account how long it takes for a commodity to be produced. Thus, it hides all the timing issues. How does one handle the launch of new commodities in a way that would not alienate certain groups (e.g. a new CPU will be released—who will get it first?). We envision the value-fairness strands of the CEP to have veto over what commodities can be created, but also employ a large bureaucracy that would try to detect what the beliefs of all parts of the population on what commodities are produced (e.g. through voting, examining consumption habits, online discussions). Finally, it would ensure that data collection methods are neither intrusive nor incomplete, allowing for the public use of collected data. In today’s terms, this strand would be a cross between an ethics board, AI ethics, animal welfare labour union board, and a consumer advocacy board.

#### 6.5 Meta-value

Institutions tend to have a self-congratulatory attitude towards their own goals. For example, if the planning board sets a goal for creating 100 houses a year, while knowing that the actual capacity is much bigger, using this number as a metric for the quality of the process, one ends up with an extremely self-referential situation where the rewards are set up in a way so as to make failure impossible. Note that obviously markets have similar problems, owing to the wealth maximisation tendencies. Ultimately, producers of goods need to be accountable to the final users of their products, and feedback needs to be constant. This part of the CEP intuition would constantly try to create new ways to measure satisfaction and improve the rest of the institution. Historical analogues of this would be “red teams” in product and policy design, as well as internal affairs units.

### 7 Final thoughts

We have discussed some ideas on how to establish and advance CEP and what its foundational principles and goals should be. We ended up with a model in which a “social factory” is gradually developed. This is in stark contrast to the current developments in manufacturing, where work is offshored and centralised (Houseman et al. 2011), with the Western workforce moving towards value-added service work. The overall model is strongly reminiscent of a “leftist” version of distributism (Pope Pius XI 1891; Chesterton 1910), a societal vision which stems from Christian thought (linked to both failures (Morris 1999) and exceptional successes (Carter 2010)). The overarching high-level vision is of a large society of independent producers, each owning the means of production of final commodities, and coordinated through an automated, online system—somewhat close to what Marx would describe as an “Association of Freely Associating Producers”. Arguably, when these ideas were initially put forward, it was impossible to implement them; information asymmetries, lack of development of productive forces, and general tendencies to create positive feedback loops resulted in primitive accumulation being reconstituted. The existence of big nationalised institutions creating the “lego” components of the economy, while a multitude of smaller players coordinate voluntarily through CEP offers a compelling alternative; at the very least, this is where the whole edifice of what we seem to know about AI points.

One might argue that the vision presented is close to classical European social democracy, and this is partially true. It does create a distinct part of the economy that would deal with commonly needed goods and services (which we have more formally defined as empowering commodities), but on the other hand, it envisions a much

more decentralised setup, with most economic activity taking place at home or at local coops.

Invariably, any plan or theory that comes in contact with reality will have to change, and what we propose here is no exception. Ideally, one would like to start small, do incremental changes, see how they behave, learn the appropriate lessons, tweak things a bit more, and so on and so forth. Without any serious experimental work, one has to rely on observational data and quasi-religious beliefs. In our case, the experimental setup would probably entail helping set up smaller settlements with the above principles in mind. Who might actually take this project (or any other project of economic transformation) forward is currently unclear.

**Acknowledgements** We acknowledge the support of the Business and Local Government Data Research Centre BLG DRG (ES/S007156/1) funded by the Economic and Social Research Council (ESRC).

**Data availability statement** No associated data.

## Declarations

**Conflict of interest** The authors have no relevant financial or non-financial interests to disclose. The authors have no conflicts of interest to declare that are relevant to the content of this article

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Adaman F, Devine P (2001) Participatory planning as a deliberative democratic process: a response to Hodgson's critique. *Econ Soc* 30(2):229–239. <https://doi.org/10.1080/03085140120042299>
- Albert M (2004) *Parecon: life after capitalism*. Verso. ISBN 9781859846988. <https://books.google.co.uk/books?id=Ij56T-acvUUC>
- Allyn-Feuer A, Sanders T (2023) Transformative agi by 2043 is <1% likely
- Anthony T, Tian Z, Barber D (2017) Thinking fast and slow with deep learning and tree search. In: Guyon I, Von Luxburg U, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R (eds) *Advances in neural information processing systems*, vol 30. Curran Associates, Inc.
- Badia AP, Piot B, Kapturowski S et al (2020) Agent57: outperforming the atari human benchmark. In: *International conference on machine learning*. PMLR, pp 507–517
- Barto A, Duff M (1993) Monte carlo matrix inversion and reinforcement learning. In: Cowan J, Tesauro G, Alspector J (eds) *Advances in neural information processing systems*, vol 6. Morgan-Kaufmann, Burlington
- Beer S (1979) *The heart of enterprise*. Classic Beer series. Wiley. ISBN 9780471275992
- Bookchin M (1982) *The ecology of freedom*. New Dimensions Foundation, Naperville
- Bortoft H (2012) *Taking appearance seriously: the dynamic way of seeing in Goethe and European thought*. Floris Books, Edinburgh
- Boyd JR (1987) *Organic design for command and control. A discourse on winning and losing*
- Carter M (2010) The landless rural workers movement and democracy in Brazil. *Latin Am Res Rev* 45:186–217
- Cartwright N, Cat J, Fleck L, Uebel TE (1996) *Otto Neurath: philosophy between science and politics*. Cambridge University Press, New York
- Castoriadis C (1974) Sur le contenu du socialisme iii: la lutte des ouvriers contre l'organisation de l'entreprise capitaliste. *L'Expérience du mouvement ouvrier, Prolétariat et organisation* 2:28–36
- Chesterton GK (1910) *What's wrong with the world*. Dodd, Mead, New York
- Cockshott WP, Cottrell A (1993) *Towards a new socialism*. Spokesman Books, Nottingham
- Curtis A, PoliticsJOE (2022) Adam Curtis on the fall of the Soviet Union's worrying parallels with modern Britain. Youtube. <https://www.youtube.com/watch?v=663vLIYBcpI>. Accessed 11 Nov 2023
- French RM (1999) Catastrophic forgetting in connectionist networks. *Trends Cogn Sci* 3(4):128–135
- Garcia J, Fernández F (2015) A comprehensive survey on safe reinforcement learning. *J Mach Learn Res* 16(1):1437–1480
- Goodhart C (2015) *Goodhart's law*, vol 227. Edward Elgar Publishing, Cheltenham
- Graeber D, Wengrow D (2021) *The dawn of everything: a new history of humanity*. Penguin, London
- Grant RM (1970) *Augustus to constantine: the thrust of the Christian movement into the Roman world*. Harper & Row, New York
- Hahnel R (2005) *Economic justice and democracy: from competition to cooperation*. Routledge, London
- Houseman S, Kurz C, Lengermann P et al (2011) Offshoring bias in us manufacturing. *J Econ Perspect* 25(2):111–32
- Kaelbling LP, Littman ML, Cassandra AR (1998) Planning and acting in partially observable stochastic domains. *Artif Intell* 101(1–2):99–134
- Klagge JC (1986) Marx's realms of 'freedom' and 'necessity'. *Can J Philos* 16(4):769–777
- Kornai J (1979) Resource-constrained versus demand-constrained systems. *Econometrica* 47(4):801–819
- Koutsoupias E, Papadimitriou C (2009) Worst-case equilibria. *Comput Sci Rev* 3(2):65–69
- Lee MSA, Floridi L, Singh J (2021) Formalising trade-offs beyond algorithmic fairness: lessons from ethical philosophy and welfare economics. *AI Ethics* 1(4):529–544
- Leontief W (1986) *Input-output economics*. Oxford University Press, Oxford
- Littman ML (1994) Markov games as a framework for multi-agent reinforcement learning. In: Cohen WH, Hirsh H (eds) *Machine learning proceedings 1994*. Morgan Kaufmann, San Francisco, CA, pp 157–163. <https://doi.org/10.1016/B978-1-55860-335-6.50027-1>
- Marx K (1875) *Critique of the Gotha programme*
- Miljković M (2017) Making automobiles in yugoslavia: fiat technology in the crvena zastava factory, 1954–1962. *J Transp Hist* 38(1):20–36

- Montias JM (1959) Planning with material balances in soviet-type economies. *Am Econ Rev* 49(5):963–985
- Morris KL (1999) Fascism and British catholic writers 1924–1939: part 1. *New Blackfriars* 80(935):32–45
- Narayanan A (2019) How to recognize AI snake oil. Arthur Miller Lecture on Science and Ethics. Massachusetts Institute of Technology
- Nieto M, Mateo JP (2020) Dynamic efficiency in a planned economy: innovation and entrepreneurship without markets. *Sci Soc* 84(1):42–66
- Nitzan J, Bichler S (2009) *Capital as power: a study of order and cre-order*. Routledge, London
- Oliehoek FA (2012) Decentralized POMDPs. In: Wiering M, van Otterlo M (eds) *Reinforcement learning: state-of-the-art*. Springer, Berlin, Heidelberg, pp 471–503. [https://doi.org/10.1007/978-3-642-27645-3\\_15](https://doi.org/10.1007/978-3-642-27645-3_15)
- Phillips L, Rozworski M (2019) *The people’s republic of walmart: how the world’s biggest corporations are laying the foundation for socialism*. Verso Books, New York
- Pope Pius XI (1891) *Rerum novarum*. Quadragesimo anno
- Roncaglia A (2009) Basic and non-basic products. Palgrave Macmillan, London, pp 60–77
- Roth AE, Erev I (1995) Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. *Games Econ Behav* 8(1):164–212
- Salge C, Glackin C, Polani D (2014) Empowerment—an introduction. In: Prokopenko M (ed) *Guided self-organization: inception*. Springer, Berlin, Heidelberg, pp 67–114. [https://doi.org/10.1007/978-3-642-53734-9\\_4](https://doi.org/10.1007/978-3-642-53734-9_4)
- Samothrakis S (2021) Artificial intelligence inspired methods for the allocation of common goods and services. *PLoS One* 16(9):e0257,399
- Schumacher EF (2011) *Small is beautiful: a study of economics as if people mattered*. Random House, New York
- Scott JC (2008) *Seeing like a state*. Yale University Press, New Haven
- Sraffa P (1960) *Production of commodities by means of commodities*, vol 1. Cambridge University Press, Cambridge
- Stark R (1996) *The rise of christianity: a sociologist reconsiders history*. Princeton University Press, Princeton
- Sutherland R (2016) Watch out for chesterton’s fences. *Spectator* 332(9811):61–62
- Sutton RS, Barto AG (2018) *Reinforcement learning: an introduction*. MIT Press, Cambridge
- Taleb NN (2007) *The black swan: the impact of the highly improbable*, vol 2. Random house, New York
- Teney D, Abbasnejad E, Kafle K, et al (2020) On the value of out-of-distribution testing: An example of goodhart’s law. In: Larochelle H, Ranzato M, Hadsell R, et al (eds) *Advances in Neural Information Processing Systems*, vol 33. Curran Associates, Inc., pp 407–417
- Thiel P, Masters B (2014) *Zero to one: notes on startups, or how to build the future*. Crown. ISBN 9780804139298
- Turner A, Smith L, Shah R et al (2021) Optimal policies tend to seek power. *Adv Neural Inf Process Syst* 34:23063–23074
- Winstanley G (1649) *The new law of righteousness*
- Yeager LB (2004) Mises and hayek on calculation and knowledge. *Friedrich A Hayek* 2(2):1
- Zubok VM (2021) *Collapse: the fall of the Soviet Union*. Yale University Press, New Haven

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.