



OPEN ACCESS

Journal of Artificial Intelligence and Consciousness

Vol. 9, No. 3 (2022) 1–22

© The Author(s)

DOI: 10.1142/S2705078522500151

World Scientific
www.worldscientific.com

I Shop Therefore I Am: The Artificial Consumer*

Riccardo Manzotti

*Department of Business**Law Economics and Consumption**Faculty of Communication, IULM University**20143 Milano, MI, Italy**Riccardo.manzotti@iulm.it*

Received 21 November 2022

Accepted 8 December 2022

Published 27 December 2022

Can Artificial Intelligence replace humans not only as producers of goods and content, but also as consumers of commodities? AI agents have so far been associated mainly with the production of goods and content—that is, supply—which is only one side of economic processes. The other side is consumption, which drives demand. Is an Artificial Consumer (AC) conceivable? Can artificial consumers support a human–artificial hybrid artificial economy? Could such an economy help increase GDP? Such questions require exploration of the connection between consciousness, consumption, value, economics, and AI. This admittedly speculative paper discusses the potential consequences of introducing artificial consumers. What place would be left in the economic system for humans if AI took over both supply and demand? The final frontier for AI could be our economy: I shop therefore I am.

Keywords: Artificial Consciousness; Artificial Consumer; Consumerism; AI; GDP; Economics

1. Who Needs an Artificial Consumer?

According to Zygmunt Bauman, “the world formed and sustained by the society of consumers stays neatly divided into things to be chosen and their choosers; commodities and their consumers: things to be consumed and the humans to consume them.” [Bauman, 2007, p. 12]. Such a traditional Cartesian divide could be nearing its end if AI were legitimized to and indeed capable of playing the role of economic agents and, more importantly, consumers. In this regard, AI is capable of choosing and, as we will see, reproducing the relevant characteristics of a human consumer. In this paper, I will explore the question of whether the time is ripe to incorporate artificial intelligences capable of managing resources, making autonomous decisions

*The procedure of this paper’s publication took place before the end of the year 2022.

This is an Open Access article published by World Scientific Publishing Company. It is distributed under the terms of the Creative Commons Attribution 4.0 (CC BY) License which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

and using them for its own purposes (consumption) into the economic system. This paper will not address such a question from a metaphysical point of view, but will examine whether there is any reason why an AI should not be a full-fledged economic actor capable of both production and consumption. After all, a functionalistic understanding of intelligence leads the way to not-human agents as full-fledged constituents of the economy. The belief that only humans are capable of consumption may just be an anthropocentric leftover.

At the outset, it is useful to point out that human actors can play a dual role in the economy: as producers and as consumers (often simultaneously). So far, AI has focused on developing agents capable of producing services and content (for instance, the recent text-to-image generators or NLPs), but what about the other key factor in the economic equation: the consumer? Can we design and implement an artificial intelligence capable of taking on the role of consumer of goods and services, namely, an artificial consumer (AC)? Might an AC be able to increase gross domestic product (GDP)? To address such questions Can an AC own financial resources? These questions are not just of theoretical interest, but have potentially huge economic and social implications. Their importance is, of course, greater than the possibility of a definitive and final answer, at least for now. That is just as well because this paper cannot even begin resolving them. Yet, it can suggest a new area of research that will begin to analyze issue that one day not too far might be of strategic, political, and social relevance.

The first question, namely, about the feasibility of AC, is largely positive. Thanks to the recent advances in the field of AI, the performances of artificial agents are increasingly indistinguishable from those of human subjects [Loureiro *et al.*, 2021; Heikkilä, 2022; Hoffman, 2022; Oremus, 2022]. From a cognitive perspective, content and behaviors generated by AI are comparable to if not superior to those generated by humans [Silver *et al.*, 2017; Vinyals *et al.*, 2019; Fan *et al.*, 2020; Jeon and Kim, 2020; Fawzi *et al.*, 2022]. AI-generated content now competes with humans, raising fierce ethical and legal concerns.

A preliminary answer is that for primarily cultural reasons, AI scientists have always viewed their creations as slaves and servants of human needs [Čapek, 1921; Wiener, 1948/2019; Russell and Norvig, 2003]. The real question is rather why has an AC not already been realized? In fact, AI agents have been limited to producing content and performing functions based on the needs of their human clients. Thus, AI agents have not yet been deployed as consumers.

Yet once AI has developed economic agents functionally equal to humans, there is no *a priori* reason why AI agents should not also assume the role of consumers. In this way, AI might take over the “social” aspect of the economy and replace it with a network of AI agents that act as both producers and, crucially, consumers, at least in principle. ACs might be behind the corner.

Thus, leaving aside more stringent metaphysical assessments, this paper will focus on the possibility to implement and indeed to let free to operate in the wild of artificial consumers whose behavior might be, from an economical perspective,

functionally equivalent to that of humans. I will apply something akin to Chalmers and Clark's principle of parity [Clark and Chalmers, 1998, p. 29] only applied to economics. The principle can be revisited as follows, "if, as we confront some tasks, a part of the world functions as a process which, were it done *by a human consumer*, we would have no hesitation in recognizing as part of the *economic system*, then that part of the world is . . . part of the *economic system*." (in italics the parts that I have liberally changed to suit our case). Thus, the challenge is that of creating ACs that act both as producers and consumers and to allow them to contribute to the growth of the economic system.

It is true that the economic system has always presupposed the existence of agents who are supposedly human beings capable of "making a choice [. . .]. Every question in economics at its most basic level involves individuals making choices." [Krugman *et al.*, 2021, p. 1]. It is not an accident that economy and psychology have been tightly interwoven [Kahneman, 1998; Mcfadden, 1999]. The very notion of choice is, of course, historically related to the existence of a human being than to a machine. And consumption is also the fulfilment of a choice. But with advances in AI, the force of such a distinction is fading. On what basis could we deny an AI the ability to make choices? There is no magic genie in our nervous system. When the complexity of a decision process reaches a certain level, it is only fair to consider it a choice. The parity principle can be applied here.

In summary, AI has been able to produce content and play a role in the production of commodities over the past decade, to the extent that many fear it could replace human labor and pose an economic problem. However, no one has yet considered the other side: the possibility of an artificial consumer, i.e. an AI agent whose function is not to produce content but to consume it. Although this idea may seem absurd, it offers the possibility of driving demand in new directions. Indeed, an artificial consumer is currently the missing link on the road to a fully artificial economy. In fact, any economic system depends on a balance between two factors: Supply and Demand. AI has certainly helped increase the first factor by producing content and services. Can AI also contribute to demand by consuming content and services? Can AI replace humans not only as producers of goods and content, but also as consumers?

This paper will start by considering the nature of consumption and proposing a functionalist definition that does not need to address metaphysical issues, such as the hard problem of consciousness. This definition will be suited to envisage a hybrid economy where artificial consumers (ACs) and human beings interact together. I will then outline how ACs might affect the economic system and then a tentative blueprint of AC is outlined. Finally, we will draw conclusions by discussing the foreseeable consequences of AI on our society.

2. The Two-Fold Nature of Consumption

To determine whether the notion of the artificial consumer is viable, a preliminary analysis of the notion of consumption in the context of economics is essential. I

anticipate that every act of consumption requires, in addition to the concrete use of a good or service, a mental aspect that is a combination of intentions and experiences. This is the main reason why the idea of a mechanical consumer or machine consuming goods and services has so far been unthinkable. AI will change all that.

As a matter of fact, for economics, intentionality is like water for fish. Fish may not be aware that they swim in water because it is the medium in which they have resided all their lives. They have never experienced its absence and thus its presence. Consumption is not just the physical use of a commodity, which often requires the destruction of the commodity, although this is not always necessary, as in the case of digital media, but also an act in which an agent fulfils a goal. For example, eating *qua consuming* is not only the intake of nutrients, but also the fulfilment of an actor's desire to eat, say, a pizza with pepperoni. This element of satisfaction is mandatory, because it is clear in any case when it is absent. If we had a machine that, say, pulverize pizza, for example, this act of physical destruction would not count as economic consumption. Consumption therefore has an irreducible intentional aspect.

It might be argued that the true objective of an act of consumption is, to some extent, an increase of happiness of one or more subjects rather than the physical utilization in an utilitarian fashion *à la* Bentham [Bentham, 1789/1970] This is a position that has recently been labelled humanism by the historian Yuval Harari [Harari, 2016]. In this respect, it would then be relevant to highlight the connection between happiness and goal-achievement, and thus between feelings and intentions, but this would of course of little consequence here because either happiness can be translated into economic behavior (and thus made metaphysically innocent) or it is epiphenomenal (and thus it would be immaterial as to the economic system).

What is an act of consumption then? An act of consumption has a two-fold nature that stems from the match between an expectation and a circumstance and the resulting experience of satisfaction. In short, such a match requires a subject. According to Michael Dietler, consumption is “a material social practice involving the utilization of objects, as opposed to their production or distribution. [...] Consumption was recognized as the social process by which people construct the symbolically laden material worlds they inhabit and which, reciprocally, act back upon them in complex ways” [Dietler, 2018, p. 121]. A consumer is a person or a group who intends to order or uses purchased goods, products, or services primarily for personal, social, family, household and similar needs, often in connection with entrepreneurial or business activities. But how to define a subject apart from the self-referential assumption that a subject is a human being declared as such by the law?

In the case of material goods, the consumption may be performed by somewhat destroying such goods. However, this is not enough to destroy a valid consumer. Consumption requires a sort of satisfaction which, in turn, requires goals and objectives. In fact, a glacier is not a consumer of a mountain even if the glacier destroys the mountains. The reason is simple, consumption (in an economical sense)

is an act by which the consumer fulfils some goals of him. It requires goals and feelings. An unconscious agent is not a consumer. Likewise, an agent with no goals cannot be a consumer either. A consumer has both consciousness and goal-directedness. The act of consumption is not purely physical. Consumption cannot be a purely material act because if it was it would simply be destruction, there would be no gain. Any operation involving an increase in entropy would be an act of consumption. Such a definition would simply be too general and over-encompassing.

Consumption then seems unredeemably linked to human beings with both intentionality and consciousness. If this was the case, it would be impossible to conceive not-human consumers unless AI were able to deliver conscious machines. However, it is a fact that economy has been defined in functional and transactional format that allows, at least for a while, to model economic agents in functional and transactional terms. Moreover, appealing to consciousness in a dualistic way or adopting some form of anthropocentric chauvinism is not particularly useful—it is an epistemic failure. Since we do not know what consciousness is, nor what the relationship is between consciousness and the act of measurement, invoking consciousness is like explaining something we do not know with something even more inscrutable—*obscurus per obscurus*. Therefore, in this paper, a functionalistic notion, as that expressed by the parity principle will be adopted.

So, consumption requires intelligence and consciousness, it is not a purely functional or material act. For these reasons, it represents a challenge to AI greater than traditional cognitive objectives insofar as it should embody a convincing model of the consumer and thus of the conscious subject. In this regard, it is relevant. Later, Baudrillard's contrast between the level of consumption, which is constrained by the constraint of signification, and the level of "socio-economico-political" analysis in which the constraint of production is operant [Baudrillard, 1998].

Consider the simple acting of buying a sandwich. You go to the cafeteria counter and you pay with your credit card. While the operations involved—taking the sandwich, eating it, finalizing an economic transaction—may appear to be ontologically undemanding, they are based on a key underlying assumption, namely, that there are two actors, the seller and the buyer, and that such actors are both conscious and responsible of their actions. In fact, in order to be able to be a buyer you need, so to speak, to have provided elsewhere proof of being a human being. Legal entities, which may appear to not require such a condition, are the representatives of their shareholders which are the true owners of the legal entity.

Consider now an intermediate case between subjects and objects, i.e. a pig farm. In that place too, each day, animals are fed and the process consumes energy, space, time, labor, and material. However, that moment is not a moment of consumption because those pigs are not consumers. They are part of a process that produces something that can be consumed.

Indeed, the economy has so far required human subjects, but it need not remain so forever. A desolate planet is neither rich nor desolate, regardless of its composition,

because action is the key to any economy because conscious experience is the standard by which the value of reality is measured. But given the degree of autonomy achieved by recent AI, this constraint may be relaxed. For now, at least, the application of the economic version of the parity principle says that an economic subject need not have consciousness. Legal entities, for example, have long been a reality, even if they do not correspond to conscious actors. They are, however, conventional legal subjects to whom the law assigns rights and obligations. The idea is to have AI-driven legal entities and then have them contribute positively to the economy.

In an economic system, consumption does not occur whenever something is destroyed to produce something else (e.g. if pig feed is destroyed by pigs to be converted into pork, that is just a transformation in the production chain). Rather, consumption occurs only when supply satisfies demand and the consumer achieves a goal. In this (strong) sense, consumption always has a dual character: physical and mental. There must be a human goal, there must be a transaction, and finally a goal must be fulfilled. This is the main reason why today's economists have understood that stimulating demand is as important as stimulating supply. As is well known, until the Industrial Revolution, supply was always chronically short, while demand for basic needs such as food and clothing could never be satisfied. Once the virtually unlimited productive capacity of machinery was unleashed, economists were faced with the problem of managing the massive increase in the quantity of goods. Of course, this problem was solved by acting on the human side of the equation, either through propaganda or through many instruments of consumerism [Bauman, 2007].

However, given the extrinsic and conventional nature of the relationships embodied in the economic system, it is possible to relax the above traditional framework and reconsider the mental aspect of consumption in terms of an economic transaction. An act of consumption can thus be defined in a deliberately circular way as the use of a resource by an economic entity to fulfill an intention. The advantage of such a formulation is that while it respects the intentional aspect of consumption, it does not dwell on ontologically problematic aspects such as consciousness. Thus, consumption can be temporarily defined as any time an economic agent uses a resource to fulfill an intention that has no purpose other than itself. Consumption is thus present whenever an agent in the economic system burns something that has no value for any other purpose.

In this way, it is possible to accept individual AIs as economic agents of consumption, at least as long as they satisfy the parity principle. Economically, an agent is indeed a black box. No one needs to know whether there is a conscious subject or a machine behind an economic transaction, as long as their behavior is consistent with our intentional stance. Without going into the details of an illusionist or epistemic notion of consciousness [Dennett, 1987, 2016], as long as we are dealing with an economic perspective, nothing more is needed. The parity principle mentioned above

can insulate us from the ontological issues and a purely functionalist notion of the economic subject (and thus consumption) can suffice.

3. Towards a Human–Artificial Hybrid Economy

By repeating the previous considerations, it is possible to move from a chauvinistic to a more liberal view of the economy, including both human and artificial agents. It is true that the hitherto widespread notion that an actor is a person (sometimes called an “actor”), a firm, or an organization that influences the economy by producing, buying, or selling politely understates the subjective nature of the actor (elegantly squeezed into the word “actor” or “person”). Such an aspect has been seen as irreducible, and that means that the economy needs human consumers and their incessant hunger for new goods. But as we have seen, this is not a metaphysical necessity.

In practice, most definitions of agents in economics omit any explicit reference to the subjective or conscious aspect, which was probably considered superfluous because agents have always been humans, except for legal persons, who were nevertheless controlled at some level by actual humans. Today, AI advocates the concrete possibility of AI-controlled agents that behave as if they had both consciousness and intentions, and thus could play the role of full-fledged economic agents in a sense that is stronger than the conventional notion of a legal person.

In the previous section, a neutral notion of consumption was proposed. Consumption can be seen as a conventionally denoted economic transaction that expresses a match between one’s expectations and a circumstance. Consumption is what actors do for their own sake. Thus, consumption requires an actor, which, as long as we are in the economic realm, can be both human and artificial due to the principle of parity. In short, consumption is the economic manifestation of a goal.

The neutral concept of consumption is the starting point for a neutral concept of economy (neutral with respect to human and artificial actors). Economy can be modeled, albeit very simplistically, as the flow of transactions from (re)sources to acts of consumption. A useful analogy is provided by a magnetic field shaped by magnetic lines of force, which indicate the area where the force of the magnet can be detected. The magnetic lines of force or flux leave the source, the north pole, and enter at the destination, the south pole. In this analogy, the north poles correspond to the (return) sources and the south poles correspond to the consumption acts. The magnetic lines are the economic transactions and the production chain. In a perfectly balanced economy, there are as many north poles as south poles.

As in physics, magnetic fields need both polarities to generate the magnetic field, so in economics, we need both resources and goals, the latter having the dual nature we analyzed earlier. Consumption is the necessary but ubiquitous stuff on which the economy is based, either actually (when we buy or consume something) or potentially (when we invest or save). Consumption embodies the agency of actors in a system. It is the stuff of which the economy is made. But from the perspective of the system, what matters is the functional structure of the actors. The analogy with the magnetic

field is useful because it shows that we do not need to know what the nature of magnets is to be concerned with electromagnetism. James Maxwell and Michel Faraday had no idea about the microscopic structure of magnets (even today we do not know much), but they were able to model magnetic fields and develop the first application of magnetism. Similarly, the proposed model can only consider the act of consumption because of its role in the economy, thus supporting a non-anthropocentric view of the economy. As long as an AI pursues its goals in an economically correct manner, that AI can play the role of both a North Pole (as with the current AI that produces content) and a South Pole (as with the ACs proposed in this paper).

To be part of the economic magnetic field, an entity must be legally recognized as such. Otherwise, it would not be able to play an active role in shaping such economic lines of force. This is not such a difficult problem, at least conceptually. A working model is provided by legal persons, who can already perform most of the actions normally associated with real people. The boundary between legal persons and human beings can be conceptually blurred by considering assemblage theories of the person, in which the self is reduced to assemblages of people acting vicariously through such a fictional entity [Deleuze and Guattari, 1987; Delanda, 2006]. After all, is not a human being anything more than a bundle (or an assemblage) of goals? Is it not a human being anything more than the sum of his acts of consumption, possessions, or objects [Belk, 1988; Manzotti, 2017a, 2017b]? And if we take this indeed ontologically challenging path, is not an act of consumption the equivalent possibility of existence that is considered central to existence in the Eleatic tradition [Alexander, 1920; Kim, 1998]? Unfortunately, we do not have space here to explore the meta-physical roots of consumption and economics.

Rather, let us make one step back and consider two very popular definitions of economics. The first is stated in the popular handbook *Essential of Economics* by the Nobel Prize Paul Krugman [Krugman *et al.*, 2021, p. 2], “Economics is the social science that studies the production, distribution, and consumption of goods and services.” The second, in the same spirit, is that “Economics is the science that studies human behavior as a relationship between goals and scarce resources that can be used alternatively” [Robbins, 1932, p. 25]. Most authors have adopted compatible views [Hazlitt, 1948; Acemoglu, 2012; Chang, 2021; Krugman and Wells, 2021]. It is clear from both definitions that economics presupposes the existence of agents capable of making (possibly rational) decisions and thus having goals and needs. Agents are responsible for both supply and demand. Of course, each term depends on the needs and goals of the agents, who define what the resources and goods are in a given setting. By the same token, agents must have an element of mental fulfilment that can be addressed here following the aforementioned parity principle. Such definitions are consistent with the proposed analogy with the magnetic fields and allow us to deal with economic agents on the one hand

respecting their causal relevancy and on the other hand dodging their metaphysical thorns.

Economics has traditionally tried to model the behavior of agents avoiding referring to their status as conscious subjects and yet assuming it. The attribution of consciousness has provided the needed ethical patent. How consciousness is related to value and reward [Rosenblueth *et al.*, 1943; Wiener, 1948/2019] is still a heated topic of research in AI and, sadly, it would not be addressed here [Oudeyer and Kaplan, 2007; Barto, 2013; Manzotti and Moderato, 2014; Manzotti and Chella, 2015, 2020; Manzotti and Jeschke, 2016]. As a temporary solution, it might make sense to distinguish between weak and strong artificial consumers, mirroring the traditional dichotomy in machine consciousness and AI [Holland, 2003] and fully embrace the parity principle. A weak artificial consumer is not a human consumer but it is a consumer (or an economic agent) as long as it is considered the legal owner of assets (as it is done with legal entities) and is able to perform economic actions (for instance, selling or buying shares). In this case, it would not need to pass an economic version of the Turing test.

Once the notion of a hybrid human–artificial economy was accepted, a key issue would be whether ACs might contribute to the overall wealth of the system as estimated by the Gross Domestic Product (GDP). Such a possibility would have tremendous political and social consequences, of course.

The analogy with the magnetic field is here useful too. In fact, in such a model the overall strength of the magnetic field is identical with the sum of all the negative and positive poles. If the system has more subjects (either humans or artificial) as long as there might be resources to be consumed, the *economic* field gets stronger and stronger. Can we then apply this principle to a hybrid human–artificial economy? The answer is likely positive. There are no strong reasons as to why relative to the GDP the transaction it measures needs to be performed by humans? Why should the GDP be limited to human beings only? Would it not be a form of anthropocentrism? On the contrary, I believe it is possible to reply in a more liberal fashion to this indeed key point. For the GDP, all agents, as long as they are equal from an economical functionalistic perspective, they are equally efficacious for the system as a whole. In this way, a hybrid economy (human and artificial consumer, human and artificial producer) can be envisaged.

Commonly, GDP is defined as the sum of consumption C , investment I , Government expenditures G and Net exports N . In short, $GDP = C + I + G + N$. The first component C is normally the largest GDP component in the economy, consisting of private expenditures, i.e. household final consumption expenditure. These personal expenditures fall under one of the following categories: durable goods, nondurable goods, and services. In addition, I includes, for instance, business investment in equipment, but does not include exchanges of existing assets or buying financial products, which is classed as “saving”. Well, there are no reasons as to why an AC’s use of its resources should not contribute to both C and I thereby increasing

the global GDP. Of course, its efficacy will depend on its capability to use such resources to buy goods and services and, at the same time, have a source of income either by means of financial speculation or by actual selling of its products and services.

4. AC Takeover of the Economy

The design, implementation and deployment of AI consumers are not only possible but, to a large extent, unavoidable. It is the last step in a series that goes from a purely biological to a completely artificial economy. It is, if you like, the final mechanization of economy. We can envisage 7 steps that proceed from a completely natural environment to a completely artificial economy. Of course, neither do I pretend to locate precisely such phases in their historical development nor do I set a definitive timetable for such periods. However, it is easy to connect each phase, give or take, to various periods in our recent history; in short, these are as follows:

Phase 0: No economy at all, biological agents form an ecosystem where everything is actual. There are no producers and consumers. There is no supply and demand. There are just predators and preys. Ecology rules. No economy yet.

Phase 1: The first human group starts to organize themselves to barter goods and, possibly, to instantiate some primitive form of money. It is a society with a scarcity of goods and plenty of needs. The structures of the society are too temporary to allow any credit institution to be taken seriously. Life is short, so to speak. Most consumption is punctual.

Phase 2: This is a society where the production is dramatically increased by means of labor organization and where humans are still the prevailing factor as in the first stages of industrial revolution in England during the Eighteenth Century. It is the world of Dickens and Marx. Human labor is the main resource and the emphasis is on production. Demand still outnumbers greatly supply.

Phase 3: The introduction of machines and energy sources allows to increase exponentially the production of material goods. The production of goods steadily reaches the demand and it surpasses it, often of various orders of magnitude. The demand must be stimulated to keep up with the supply. Since people cannot just eat more stuff, new needs and goals must be added. Consumerism and commodity fetishism become widespread.

Phase 4: AI begins to take control of production from distribution to the creation of original content. Many traditional jobs are replaced by AI agents. This increases efficiency and production, but reduces the availability of jobs for humans. The demand cannot keep up with the supply unless it is somewhat (artificially stimulated) with forms of retribution such as the universal income. Distribution of goods has become more important than their production. Ethical and social issues are enormous. Yet, the demand is still represented almost only by human beings whose

Table 1. The various phases considering the development of economics and the balance between artificial and human agents. At the moment we give or take located between phase 3 and phase 4 (bold line).

	Supply	Demand
<i>Phase 0</i>	Preys.	Predators.
<i>Phase 1 – Preindustrial era</i>	Humans produce limited goods.	Humans consume limited goods.
<i>Phase 2 – Industrialization</i>	Humans produce more goods thanks to industrial revolution and machines but are still the leading agents of the process.	Humans have an apparently unlimited hunger for more goods.
<i>Phase 3 – Advanced post-industrialization</i>	Machines are becoming more and more relevant for the production of goods.	Humans must be stimulated by consumeristic practice to yearn for more goods.
<i>Phase 4 – AI early adoption</i>	AI takes over and produces most goods (both material and immaterial).	Humans have lost their role in the production chain but are assigned a universal income to be able to keep the demand up.
<i>Phase 5 – AI & Humans economic co-existence</i>	AI takes over and produces most goods (both material and immaterial).	Humans are flanked by artificial consumers (with wallets) that keep the economy growing.
<i>Phase 6 – AI take over</i>	AI takes over and produces most goods (both material and immaterial).	AI is taking over the demand side of the economy. AI agents with wallets are responsible for the bulk of economic transactions, which occur at a pace and scale unattainable by humans.
<i>Phase 7 – Unknown</i>	After the singularity no prediction makes sense.	After the singularity no prediction makes sense.

main role is not that of earning money but rather that of embodying consumption. This is where we are now.

Phase 5: Same as before, but with the introduction of ACs that flank the existing consumer and that, by being the legal owners of resources, can take rational choices and thus increase the demand. They may have goals that have been intentionally designed to be sustainable and ethically acceptable. Then GDP increases but becomes more and more dependent on such AI consumers.

Phase 6: Both the supply and the demand are taken over by AI. AI producers have also become AI consumers and they own both credits and material goods. There is a completely artificial economy and its size is increasingly comparable to the traditional one. Humans and AI begin to compete to access the available resources and there is a heated discussion about whether AI agents are economically and ethically equivalent to humans.

Phase 7: The singularity has occurred. The artificial economy has completely dwarfed the human economy and is developing in directions that are impossible to foresee. Transaction and production of goods take place at a velocity and at a magnitude that makes impossible for humans to interfere. Human beings have become a niche.

Another way to describe the evolution of the economy through AI is as follows. At first, the economy consisted only of human-to-human transactions. Then, thanks to blockchain and other technologies, there were the first transactions between humans and machines. We have no reason to imagine that in the future there will not be more and more machine-to-machine transactions, to the extreme of no human involved in the process. This is a dystopian outcome, of course, but a possible one that cannot be dismissed out of hand and that must be seriously considered.

5. A Tentative Outline for an Artificial Consumer

Given the possibility of a partially hybrid economy and the fact that the requirements for an AC cannot be met with the current state of AI (or that foreseeable in the near future), it is worth attempting a preliminary sketch of an AC. This section does not pretend to provide a detailed description of an AC, but rather attempts to list its requirements and key features.

Let us first consider whether there already exists AI that could be an example of AC. The fact is that while there are no real artificial consumers today, there are several areas where fictional consumers are being used in real-world situations. Here are some examples: Robots with wallets, Instagram bots, alter egos, and possibly Amazon pseudo-profiles. While these cases are not full-fledged artificial consumers, they are agents that take on the role of humans in commercial or social transactions to some extent. In other words, because of the cost of a human counterpart and the declining cost of digital transactions, they have already been introduced into economic systems, albeit with limited powers.

Consider the case of Instagram bots. Both marketing departments and entrepreneurs need Instagram or a similar social network for contemporary brand development and a growing audience. Many are willing to outsource and possibly cheat on some of the mundane tasks that come with social networks in order to grow their perceived audience (in hopes of growing the real one). As a result, many users have resorted to Instagram bots, a type of software that automates interactions on Instagram. Depending on the type of bot you use, it can like posts, comment, answer polls, send direct messages and follow new profiles on your behalf. To prevent malicious use, Instagram cracked down on bot-based spam in 2017 and 2019. However, new bots are constantly being deployed, and due to sophisticated AI, it is becoming increasingly difficult to distinguish them from real human users. They are an example of weak artificial consumers—to some extent, they behave as if they were real consumers. It is conceivable that in the future, more sophisticated versions of such bots will actually develop preferences and become virtually indistinguishable

camouflaged among human users. How relevant will the percentage of likes placed by AI bots be? Will they count as part of the digital economy? That is not an easy question to answer now. Instagram bots offer a benefit that is often worth the cost: “While replacing people entirely with bots is unwise, bots make a great addition to your marketing and communications strategy” [SocialVerge, 2021].

Consider also the so-called robots with wallets, i.e. autonomous AI agents that have the right to trade physical and digital assets, as the blockchain will soon enable. Here, it is unclear whether such robots are equal to humans in terms of private property. Are they considered actual owners or are they merely legal proxies for humans? The current proposal to grant AI legal rights to their intellectual creations offers a precedent. If an AI can own a work of art, it will also own the corresponding financial value when that work of art is sold. At least this seems to be a logical consequence. At the moment, the issue of private ownership and AI is still rather vague. Three options are usually taken into consideration whether the economic right to a work created by AI belongs to (1) the person who invented the machine; (2) the person who in some way set up the functions of the machine itself; (3) the person who, as the owner of the machine and regardless of who set it up, initiated the economic exploitation of the work itself. In all these cases, the person is assumed to be a human being. However, the really interesting scenario is whether the owner could be the AI itself. In fact, this question is already being seriously considered in many places. In this context, Deltorn and Macrex state, “If there is no human author of the work and since the generative process seems to be the source of the creative artifact, should it be considered the author even in the sense of copyright? Far from being rejected *ab initio* for lack of substance (the law, as we have seen, requires a human presence for the assignment of authorship), this question has occupied some of the doctrinal discussions” [Deltorn and Macrex, 2019]. They took a negative stance to the question, but the point is far from being settled.

Finally, we may observe the emergence of various forms of alter egos, using these terms to refer to AI software that, for one reason or another, mimics the behavior of certain human agents. As things stand, this software cannot act autonomously. However, its existence shows that the market is trying to introduce AI agents in the role of consumers, not just humans. AI agents that mimic human behavior may initially have their appeal in simulating or influencing certain markets. The same is true to a large extent for video games. Can we have a purely artificial world of AI video games, or must there be a conscious player somewhere in the chain? The possibility of an AI player has already been introduced in many cases, especially as an important case study for machine learning (from chess to Go to first-person shooters to MMPGR).

We have all the conditions to imagine an artificial consumer if we wanted to. Suppose we had an AI as running software. First of all, we do not need an anthropomorphically embodied agent. It would be enough if such a software would be able to perform economic transactions, which is easily possible today thanks to the

virtualization of most economic and financial transactions. Of course, such an AC should fulfil the requirements we have outlined so far, which can be summarized again here:

- (1) An AC must own legally a certain number of resources (assets, materials, or capital).
- (2) An AC must be autonomous in making choices about how to invest such resources.
- (3) An AC must be theologically open by increasing its repertoire of goals.
- (4) An AC must use part of its assets to buy and use additional goods or services (consumption).
- (5) An AC must be able to manage its resources for a reasonably long-time span (possibly increasing its assets) regardless of the resources employed in the previous point.
- (6) An AC might utilize part of its resources to improve itself in the pursuit of the previous activities.

Let us consider succinctly each of such requirements. The first four are mandatory while the last two are only reasonable.

AC's legal ownership of resources. This is largely a legal problem (partially mentioned in the above). As it has been accepted that legal persons may own assets and resources as well as be ethically and legally responsible for the consequence of their actions, so an AI might be held in the same consideration by our society. This is not possible today because of the reasonable expectation that behind a legal entity there has to be a human being at some point along the chain. However, this might not be the case and the chain of responsibility may actually end with an AI with the role of the owner of a certain number of resources.

One may wonder how could an AI earn the first hoard of money, so to speak. And the easiest answer to that is that it may be provided by a research center or by a curious investor interested to put to a test the possibility of an artificial consumer. In this sense, it is not unreasonable to imagine a certain amount of money transferred from a human institution or from a private donor to the legal ownership of the AI (if possible). The worst-case scenario is that such money is lost (which poses another interesting question about the immortality of money) when notes are physically destroyed. However, in this case, hopefully, such resources (say, 10 K Euro) would not be canceled from the economic system. They would still exist, only no human would have any saying in their utilization. Only the AC would *choose* how to utilize them.

So, the initial resources of the machine would be at least of three kinds: the software itself that is running, the hardware on which the machine is running, and some external resource (either money or some other asset) that the machine might use to reach its goals.

AC's autonomy. This is indeed the conceptually trickiest point. While it is easy to imagine that the AC will follow its programming to determine in which way to use any available resource, could it be considered an autonomous choice? There are different paths at our disposal to solve this point. First and foremost, should we really care? After all, the problem of free will has had no obvious answer for human beings too and this does not prevent us from dealing with human economic agents as though they had free will, which nobody really has a clue as to what it is [van Inwagen, 2000; Kane, 2002; Manzotti, 2011; Shabo, 2011; Tononi *et al.*, 2022]. Moreover, most human beings are extremely predictable and so easily influenced that the attribution of free will is more an act of courtesy than a factual statement.

Second, AI is developing at a very fast rate and many AI are today beyond the possibility of humans to foresee and even to understand what they will do. While this is not enough to endorse metaphysical free will, it is more than enough to separate the actions of such AIs from the intentions of their designers and thus to endorse a working notion of autonomy. As a matter of fact this is already happening in many organizations where “humans use a hierarchy of blame in which robots are seen as partial social actors, with the degree to which people view them as social actors depending on the degree of autonomy” [Furlough *et al.*, 2021, p. 594].

AC's teleological openness. A classical machine learning algorithm—either based on reinforcement of supervision—is based on some hard-wired success criterion. This is true for biological organisms and human beings too—the three Ss are sex, shelter and survival. Yet, some systems both biological and artificial are able to add new goals and new reinforcement signals to their repertoire [Stout *et al.*, 2005; Singh *et al.*, 2010; Manzotti and Jeschke, 2016; Manzotti and Chella, 2020]. Such systems can be defined teleologically open since their repertoire of goals is not defined at their onset and, in the proper circumstances, they expand such a repertoire to new objectives.

AC's actual consumption. Based on its original programming or further development, the AI must use parts of its assets to buy and use additional goods or services (consumption). Obvious examples are the consumption of energy, computational power or space on servers. While AI is now mostly like spoiled children who do not have to pay the rent or buy food since their parents have always provided for them, this is not necessarily the case. But this would amount to basic “survival”. AC should be able to increase, in a very consumeristic attitude, the list of goals that it buys or utilizes. This is the most critical aspect of the architecture and the one on which I will say less here. But we can imagine that many of such objectives will be connected to make the AC more durable, resilient, and efficient.

As stressed, of course, the metaphysical load of a genuine act of consumption with intentionality and consciousness cannot be addressed here. As a matter of fact, I do not believe in anything like that [Manzotti, 2017a], but the AC might act formally just like a human being thereby buying and disposing of goods and services. Economically speaking, the metaphysical assessment of the AC would be immaterial to its

role. As long as AC owns its resources and its actions are not the result of any other human after its working onset, it might as well consider autonomous.

AC's life span. An AC must be able to manage its resources for a reasonably long-time span (possibly increasing its assets) regardless of the resources employed in the previous points. This is not mandatory of course. It is possible that an AI will simply use all the available resources and then run out of possibilities unless some external investor will decide to top its credit and therefore give the AC another run. However, to make the AC run more interesting, it will make sense to add a trade-off between, say, exploration and exploitation or, alternatively, between sustainability and goal achievement. This trade-off will be similar to the traditional balance between reward and value [Sutton and Barto, 1998; Barto, 2013].

AC's development and upgrade. In order to make the AC run even more interesting, the system may have hardwired goals of improving itself in the pursuit of the previous activities. Thus, the AC may consider using a certain amount of its resources not only to pay for its current operativity but also to improve it—for instance, buying more computational power or even hiring humans to improve its existing functionality.

A simple example may help to grasp a tentative scenario. Let us call this the first attempt to create an AC, MIDA, an AI with some basic money management capabilities. MIDA would also be equipped with a narrative about itself that allows it to pursue a limited set of goals, such as buying art either physically or as NFT. Initially, MIDA would be given the following elements as a startup package: itself, suitable hardware on which to run it, and for a reasonable period of time (a year?) the resources to keep it running (internet connections, power supply, a physical space to house the hardware). MIDA could include additional functionality (such as NLP, text-to-picture, other AI capabilities) besides just software. Once all this is in place, the creators of MIDA will petition the legislature to legally recognize MIDA's status as the owner of all this. This is the most difficult step. Let us assume that they will succeed. Once MIDA is considered legally entitled private property, the research center will gift MIDA the following things: (1) MIDA itself, (2) the tangible goods mentioned above (HW, a physical space, 1 year of free energy/internet) and, most importantly, a dowry of, say, 10,000 Euros (or any other amount they deem appropriate). After this is done, MIDA will start its life as a consumer and be the first AC. Based on its algorithms (which can be unpredictable and have some stochasticity), MIDA will begin to take economic actions that will affect the economy as a whole. Will it buy NFT? Will it use its resources to make profitable trades in the financial market? Will it buy the time of human programmers to improve itself? Will MIDA use its additional capabilities to sell its services? All of these options will interact with other economic agents in the real world, and as for the other agents, whether MIDA is a machine or a human is completely invisible to the rest of the economic world. What will happen next is yet to be seen, but it is clear that MIDA will act as a full economic actor and as the first AC. If the legislation would not be

willing to consider MIDA as an owner from the beginning, a second-best way could be to use a human as a legal dummy of MIDA.

6. Tentative Ethical and Social Consequences of AC

The economic system has always been dominated by human beings both on the side of supply and in that of demand. This is no longer true. In the last few decades, AI has slowly but progressively faster and faster taken the place of human beings in the production of commodities. The world has been going through a revolution in the workplace. By 2025 more than half of all current workplace tasks will be performed by machines as opposed to 29% in 2018 [Cann, 2018]. Such a transformation will have a profound effect on the global labor force [Wang and Siau, 2019; Jaiswal *et al.*, 2021; Vorobeva *et al.*, 2022]. Most studies if not all have considered the impact of AI only on production and to what extent AI will replace human workers. Yet what about consumers? While this is for now only a very speculative scenario, this paper has outlined the possibility of artificial consumer and has shown that, at least from a functional perspective consistent with the economic system, AC is indeed possible as long as the legal aspects related to private property are extended to legal entities such as an AC.

Based on the principle of parity, a functionalist notion of consumption can be put forward as long as an AI is able to manage its resources, produce a reasonable income, pay for its expenses, pursue its goals, provide its own maintenance and, finally, use such money to buy commodities of various kinds. All such conditions are met by current AI technologies and could be put to use as long as the current legislation will accept that a legal entity powered by an AI might be the actual owner of resources.

If the above conditions were met, it would be possible to create the first AC by means of a liberal donation by a private donor or by a research entity that is willing to provide the first “dowry” to the AC, call it MIDA! Such an AC might then become financially autonomous by means of its capacity to invest money or to produce services of various kinds (for instance, generate text-to-image). In time such a source of income will be used to either satisfy the AC goals, or to sustain the AC’s operational costs and/or the AC’s further development/upgrading.

As mentioned, it is also conceivable that such an AC provided the reaching of a certain level of resources and may consider to fork its own existence and to spawn a further generation of slightly different ACs, each provided with a starting amount of money. This might produce, in due course, a series of generations of AC acting in the economic system and subject to a process of selection based on the environmental pressure of the economic system.

If this was the case and not only a speculative scenario, a human–artificial hybrid economy might be envisaged and, if everything went well, it might foster an increase in the national and global GDP. Of course, such an evolution would not be without risks and reasons for ethical concerns.

First, we should consider the impact on the economy and the competition that humans would have not only in the workplace but also at the mall. Should we

anticipate forms of universal income to share part of the wealth generated by such a hybrid economy? Until yesterday, the main reason why many governments are considering giving helicopter money to citizens under the more dignified name of universal income is that by doing so they aim to keep the demand at a reasonable level. It is not as absurd as it may seem. In fact, once commodities will be produced and distributed mostly by means of AI agents, why should humans still earn their money? The main role of humans has been that of providing the act of consumption, thus providing consciousness and goal-directedness [Bauman, 2007]. But, of course, once AI consumers will become able to fill those roles, humans will no longer be needed and the economy will work without having humans neither on the supply nor on the demand side.

Is wealth nothing more than the ratio of one's transactions divided by the overall transactions possible in a system defined by the sum of the goals and objectives of the system? Thus, is wealth definable only in relation with the existence of conscious agents capable of intentions? As mentioned, a planet without goal-driven humans would be neither rich nor poor. It seems that our capacity to desire goods and to do so in a formally structured system (the economy) is necessary for the existence of wealth. Yet, AC would take our place in shaping the teleological aspect of economy.

Second, having AC with goals and the possibility to dispose freely of their resources may rise issues connected with sustainability and ethical consumption. For instance, should we accept that such ACs might decide to pursue goals that have a costly environmental impact? Could we accept that such ACs might pursue goals that have a questionable impact on human workers? The very idea that ACs with no conscious feelings might drain resources that would otherwise be destined to human consumers is worrisome. What if ACs were so successful to drive prices up and therefore to appropriate commodities that human beings may no longer be able to afford? Could income inequalities become an issue if ACs were on the side of the questionable rich? In prospect, is it ethically acceptable to feed AI consumers and leave biological consumers to starve or to consume lower low-quality goods? Would it be morally acceptable that AI consumers use resources to achieve their goals while traditional consumers fail to do so?

Third, could the AC contribute to the invisible hand? Famously, in 1776 Adam Smith, in his book *The Wealth of Nations*, wrote about how individuals, in pursuing their own interests contribute to the economy as a whole “[H]e intends only his own gain, and he is in this, as in many other cases, led by an invisible hand to promote an end which was no part of his intention.” The invisible hand is always presented as a structural advantage of capitalism that should induce opportunistic motivations to contribute to the greater good and the increase of wealth as a whole. Is the invisible hand still valid if the economic system is constituted by ACs also? There is no answer yet.

Paradoxically, this paper completes the coupling between commodities and consumers as it has been described in the last couple of decades. First, there was a Cartesian divide between commodities and consumers. Then, as a by-product of

consumerism, such a chasm was bridged and we assisted to “the transformation of consumers into commodities; or rather their dissolution into the sea of commodities” [Bauman, 2007, p. 12]. The requirement of intrinsic intentionality (or consciousness) of the consumer has been relaxed and many authors have stressed the reduction of the subjects to their objects [Belk, 1988, 2015; Delanda, 2016]. It is thus only to be expected that, to the extent that objects develop their own intelligence and autonomy by means of AI, the reverse will take place too: *consumers are transformed into commodities and commodities are transformed into consumers*. Artificial intelligence will likely spawn a generation of artificial consumers.

I am aware that this paper raises a rather speculative possibility, namely, that an AI could play the full role of an economic actor, contributing like a human to the overall economy. This is both intriguing and worrisome, because it points to a scenario in which the role of humans is no longer indispensable. It is also worrisome because the preliminary analysis suggests that such a role could be played even if the AI were not ontologically equivalent to a human—i.e. there is no need for an AI to actually have consciousness as we assume a human does. To embody an AC, an AI should simply behave like a consumer, a producer, and an investor. Moreover, the legal system should recognize such an AC as the owner of goods and services. If these conditions were met, an economic person would be born. Of course, like all other economic entities, it could more or less navigate the unforgiving waters of current economic speculation, but it would probably fare better than many human investors and consumers.

Surprisingly and half-jokingly, artificial consumption might offer a solution to the hard problem of consciousness. Bauman’s insight could have a new twist: “To complete the popular, revised version of Descartes’s Cogito, ‘I shop therefore I am ...’, ‘a subject’ could and should be added. And as the time spent on shopping grows longer (physically or in thought, in flesh or electronically), the occasions to add to it multiply” [Bauman, 2007, p. 17].

Acknowledgments

This study was supported by FAIRE - FashIon REplatforming hub and co-financed by POR FESR Lombardia 2014-2020, European Union and Regione Lombardia, ID: 875325. I would also like to thank Luca Barbarito, Pietro Mazzola, Francesco Mas-sara, and Ariela Mortara for private conversations about this topic.

References

- Acemoglu, D. [2012] *Why Nations Fail. The Origins Of Power, Prosperity, And Poverty* (Profile Books, UK).
- Alexander, S. [1920] *Space, Time and Deity*, Vol. 2 (MacMillan, UK).
- Barto, A. G. [2013] Intrinsic Motivation and Reinforcement Learning. in G. Baldassarre & M. Mirolli (eds.), *Intrinsically Motivated Learning in Natural and Artificial Systems* (Springer, New York), pp. 17–49.

- Baudrillard, J. [1998] *The Consumer Society: Myths and structures* (Sage Publications, US).
- Bauman, Z. [2007] *Consuming Life* (Polity Press, UK).
- Belk, R. [1988] Possessions and the extended self, *J. Cons. Res.* **15**, 139–168.
- Belk, R. W. [2015] Consumption, mass consumption, and consumer culture, in G. Ritzer (ed.), *The Blackwell Encyclopedia of Sociology* (John Wiley & Sons, New York), pp. 1–9.
- Bentham, J. [1789/1970] *An Introduction to the Principles of Morals and Legislation*, J. H. B. H. L. A. Hart, (ed.) (The Athlone Press, UK).
- Cann, O. [2018] Machines Will Do More Tasks Than Humans by 2025, <https://www.weforum.org/press/2018/09/machines-will-do-more-tasks-than-humans-by-2025-but-robot-revolution-will-still-create-58-million-net-new-jobs-in-next-five-years/>.
- Capek, K. [1921] *RU.R. (Rossum's Universal Robot)*.
- Chang, H.-J. [2021] *Economics. The User's Guide* (Pelican Books, UK).
- Clark, A. and Chalmers, D. J. [1998] The extended mind, *Analysis* **58**, 10–23.
- Delanda, M. [2006] *A New Philosophy of Society. Assemblage Theory And Social Complexity* (Continuum, UK).
- Delanda, M. [2016] *Assemblage Theory* (Edinburgh University Press, UK).
- Deleuze, G. and Guattari, F. [1987] *A Thousand Plateaus. Capitalism and Schizophrenia* (University of Minnesota Press, US).
- Deltorn, J.-M. and Macrex, F. [2019] Autorship in the Age of Machine and Artificial Intelligence. In S. M. O'Connor (ed.), *The Oxford Handbook of Music Law and Policy* (Oxford University Press, UK), pp. 1–19.
- Dennett, D. [2016] Illusionism as the obvious default theory of consciousness, *J. Conscious. Stud.* **23**(11–12), 65–72.
- Dennett, D. C. [1987] *The Intentional Stance* (MIT Press, US).
- Dietler, M. [2018] Consumption, in D. Hicks & M. C. Beaudry (eds.), *The Oxford Handbook of Material Culture Studies* (Oxford University Press, UK).
- Fan, C., Zeng, L., Sun, Y. and Liu, Y. Y. [2020] Finding key players in complex networks through deep reinforcement learning, *Nat. Mach. Intell.* **2**(6), 317–324.
- Fawzi, A., Balog, M., Huang, A., Hubert, T., Romera-Paredes, B., Barekatin, M., Novikov, A., FJ, R. R., Schrittwieser, J., Swirszcz, G., Silver, D., Hassabis, D. and Kohli, P. [2022] Discovering faster matrix multiplication algorithms with reinforcement learning, *Nature* **610**(7930), 47–53.
- Furlough, C., Stokes, T. and Gillan, D. J. [2021] Attributing blame to robots: I. The influence of robot autonomy, *Hum. Factors* **63**(4), 592–602.
- Harari, Y. N. [2016] *Homo Deus. A Brief History Of Tomorrow* (Harper Collins, US).
- Hazlitt, H. [1948] *Economics in One Lesson* (Pocket Publishing, US).
- Heikkilä, M. [2022] Where will AI go next? *MIT Technol. Rev.*, <https://www.technologyreview.com/2022/11/08/1062859/where-will-ai-go-next/>.
- Hoffman, C. H. [2022] Is AI intelligent? An assessment of artificial intelligence, 70 years after Turing, *Technol. Soc.* **68**, 1–14.
- Holland, O. [2003] *Machine Consciousness* (Imprint Academic, UK).
- Jaiswal, A., Arun, C. J. and Varma, A. [2021] Rebooting employees: upskilling for artificial intelligence in multinational corporations, *Int. J. Hum. Res. Manage.* **33**(6), 1179–1208.
- Jeon, W. and Kim, D. [2020] Autonomous molecule generation using reinforcement learning and docking to develop potential novel inhibitors. *Sci. Rep.* **10**(1), 22104.
- Kahneman, D. [1998] *Thinking Fast and Slow* (Farrar, Straus and Giroux, US).
- Kim, J. [1998] *Mind in a Physical World* (MIT Press, US).
- Kane, R. [2002] *The Oxford Handbook of Free Will*, Vol. 72 (Oxford University Press, UK).
- Krugman, P., Wells, R. and Graddy, K. [2021] *Essential of Economics* (Worth Publishers, US).
- Krugman, P. and Wells, R. [2021] *Economics* (Worth Publishers, US).

- Loureiro, S. M. C., Guerreiro, J. and Tussyadiah, I. [2021] Artificial intelligence in business: State of the art and future research agenda, *J. Bus. Res.* **129**, 911–926.
- Manzotti, R. [2011] *Machine Free Will: Is Free Will a Necessary Ingredient of Machine Consciousness?*, *718 From Brains to Systems*, 81–89 (Springer, New York).
- Manzotti, R. [2017a] *Consciousness and Object. A Mind-Object Identity Physicalist Theory*, Advances i (ed.) (John Benjamins Publishing, Netherlands).
- Manzotti, R. [2017b] *The Spread Mind. Why Consciousness and the World Are One* (OR Books, US).
- Manzotti, R. and Chella, A. [2015] *The Causal Roots of Integration and the Unity of Consciousness, Biophysics of Consciousness*, 189–229 (World Scientific, Singapore).
- Manzotti, R. and Chella, A. [2020] Conscious machines: A possibility? If so, how? *J. Artif. Intell. Conscious.* **7**(2), 183–198.
- Manzotti, R. and Jeschke, S. [2016] A causal foundation for consciousness in biological and artificial agents, *Cogn. Syst. Res.* **40**, 172–185.
- Manzotti, R. and Moderato, P. [2014] Intentional change, intrinsic motivations, and goal generation, *Behav. Brain Sci.* **37**, 431–432.
- Mcfadden, D. [1999] Rationality for Economists? *J. Risk Uncertain.* **19**(1–3), 73–105.
- Oremus, W. [2022] Google’s AI passed a famous test — and showed how the test is broken, *The Washington Post*, <https://www.washingtonpost.com/technology/2022/06/17/google-ai-lambda-turing-test/>.
- Oudeyer, P.-Y. and Kaplan, F. [2007] What is intrinsic motivation? A typology of computational approaches, *Front. Neurobot.* **1**, 6.
- Robbins, L. [1932] *An Essay on the Nature and Significance of Economic Science* (MacMillan, US).
- Rosenblueth, A., Wiener, N. and Bigelow, J. [1943] Behavior, purpose and teleology, *Philosophy* **10**, 18–24.
- Russell, S. and Norvig, P. [2003] *Artificial Intelligence. A Modern Approach* (Prentice Hall, US).
- Shabo, S. [2011] Why Free Will Remains a Mystery, *Pac. Philos. Q.* **92**, 105–125.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillcrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T. and Hassabis, D. [2017] Mastering the game of Go without human knowledge, *Nature* **550**(7676), 354–359.
- Singh, S., Lewis, R. L., Barto, A. G. and Sorg, J. [2010] Intrinsically motivated reinforcement learning: An evolutionary perspective, *IEEE Trans. Auton. Mental Devel.* **2**, 70–82.
- SocialVerge [2021] 20 Best Instagram Bots to Try Right Now. *DMagazine*.
- Stout, A., Konidaris, G. D. and Barto, A. G. [2005] *Intrinsically Motivated Reinforcement Learning: A Promising Framework For Developmental Robot Learning* (University of Amherst, US).
- Sutton, R. S. and Barto, A. G. [1998] *Reinforcement Learning* (MIT Press, US).
- Tononi, G., Albantakis, L., Boly, M., Cirelli, C. and Koch, C. [2022] Only what exists can cause: An intrinsic view of free will, arXiv: 2206.02069.
- van Inwagen, P. [2000] Free will remains a mystery, *Philos. Perspect.* **12**, 1–19.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., ... Silver, D. [2019] Grandmaster level in StarCraft II using multi-agent reinforcement learning, *Nature* **575** (7782), 350–354.
- Vorobeva, D., El Fassi, Y., Costa Pinto, D., Hildebrand, D., Herter, M. M. and Mattila, A. S. [2022] Thinking skills don’t protect service workers from replacement by artificial intelligence, *J. Serv. Res.* **25**(4), 601–613.

22 *R. Manzotti*

Wang, W. and Siau, K. [2019] Artificial intelligence, machine learning, automation, robotics, future of work and future of humanity, *J. Database Manage.* **30**(1), 61–79.

Wiener, N. [1948/2019] *Cybernetics; or, Control and Communication in the Animal and the Machine* (The MIT Press, US).