Electronic Markets (2023) 33:36
https://doi.org/10.1007/s12525-023-00649-0

FUNDAMENTALS

Machine economies

Eduard Hartwich[1] · Alexander Rieger[1] · Johannes Sedlmeir[1] · Dominik Jurek[2] · Gilbert Fridgen[1]

Received: 14 October 2022 / Accepted: 29 March 2023 / Published online: 22 July 2023

Abstract

This fundamentals article discusses efficient machine economies in which non-human agents can autonomously exchange information and value. We first identify criteria for achieving Pareto efficiency in such economies by drawing on the Coase Theorem. We then translate these economic criteria to technical requirements before developing a framework that characterizes four types of machine economies. We discuss real-life examples for each type to highlight key challenges in achieving Pareto efficiency. In particular, we highlight that machine economies with human involvement in economic interactions and governance face significant challenges regarding perfect information, rationality, and transaction costs. Machine economies without human involvement, in turn, promise a high degree of Pareto efficiency, but there are still many open questions, particularly regarding machine-enforced governance. We conclude with opportunities for future research on the interactions and governance in machine economies.

Keywords Algorithmic governance · Autonomous agents · Digital platform · Economic efficiency · Machine economy · Transaction costs

JEL Classification D01 · D40 · D49 · D50 · D52 · D60 · D80 · D81

Introduction

Digital technologies continuously evolve, transform, and merge to create innovative ways of economic interaction, not only between machines and humans but also among machines themselves. As a result, "interconnected machines, software and [digital] processes" (Arthur, 2017, p. 3) are increasingly facilitating and shifting value exchange into virtual economies. Algorithmic trading is one of many examples. It relies on software agents that autonomously observe market movements, automatically make decisions, and submit and execute orders. In effect, these software agents are fully-fledged market participants. In many instances, algorithmic trading agents have become so relevant that they account for most of the trading volume and liquidity provided on several exchanges (Hendershott et al., 2021; Moriyasu et al., 2018).

These developments are not exclusive to financial services. Autonomous agents also play an important role as value creators and contributors on digital platforms (Hein et al., 2020). For instance, the Amazon AWS IoT platform allows machines to share wear and tear data and to automatically order new parts (Amazon Web Services, 2022). Moreover, the recent improvements in artificial intelligence may lead to an increasing number of business decisions being made by software agents with little or no human oversight (Berente et al., 2021). In these and many other cases, machines engage in economic interactions, creating what can be described as a *machine economy*.

Drawing on current discussions in academia and practice, we define machine economies *as economic systems*

Responsible Editor: Christian Matt

✉ Eduard Hartwich
eduard.hartwich@uni.lu

Alexander Rieger
alexander.rieger@uni.lu

Johannes Sedlmeir
johannes.sedlmeir@uni.lu

Dominik Jurek
dominik_jurek@berkeley.edu

Gilbert Fridgen
gilbert.fridgen@uni.lu

1 SnT - Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg, Luxembourg, Luxembourg

2 Haas School of Business, University of California, Berkeley, California, USA

that involve economically autonomous machines. So far, discussions on machine economies have focused primarily on their technical building blocks. Examples include technologies that enable machines to act autonomously (e.g., artificial intelligence), technologies that facilitate value exchange with a lower degree of human governance (e.g., distributed ledger technology), and infrastructures that enable modular architectures upon which autonomous business processes can be implemented (e.g., digital platforms) (Bons et al., 2020; Hein et al., 2020; Jöhnk et al., 2021; Schreieck et al., 2016). There is also research dedicated to the economic aspects of each of these technologies (Davidson et al., 2016) and their impact on human economies (Brynjolfsson et al., 2019). However, few overarching considerations exist regarding the contributions these technologies could make to economic systems with machine participation (Lee et al., 2010) and how these systems can be governed (Leiding et al., 2021).

In particular, it remains unclear whether machine economies have the potential to avert failures that prevent human economies from reaching Pareto efficiency (Mazzucato, 2016). In this article, we thus explore the concept of *efficient* machine economies. First, we discuss sufficient criteria for such economies, drawing on classical and neoclassical economic theories, before translating them into technical requirements for machine participants. In the second step, we explore challenges related to fulfilling these criteria and requirements in four different types of machine economies. We close by indicating where future research may be required to address these challenges and facilitate more efficient machine economies.

Economic criteria for a Pareto efficient economy

Studying how markets and economies can achieve efficiency is a centuries-old endeavor. One of the earliest notions of these efficiency considerations is Adam Smith's "invisible hand", which builds on the idea that individuals acting in their own self-interest will produce greater overall benefit for society (Blaug, 2007; Makowski & Ostroy, 2001). Arrow and Debreu formalized these considerations in what later became known as the two Welfare Theorems (Arrow, 1951, 1954), which examine how and under which circumstances markets can achieve Pareto efficiency (Stiglitz, 1991).

The first Welfare Theorem states that in economic equilibrium – that is, a state where supply and demand are equal (Arrow, 1966) – complete markets in perfect competition will achieve Pareto efficiency (Arrow, 1951) and maximize economic welfare (Hicks, 1941). Pareto efficiency is defined as a situation where the allocation of resources cannot be changed in a way that makes one party better off without making another one worse off (Stiglitz, 1981). The first Welfare Theorem sees complete markets in perfect competition as a sufficient requirement for this form of efficiency. Markets are complete when "there is a market for every good" (Flood et al., 1991, p. 32) and when every agent can exchange with every other agent at negligible transaction costs (Buckle & Thompson, 2020). To achieve perfect competition, the economic literature commonly assumes requirements such as: no externalities (a situation where actions positively or negatively affect a third party) (Arrow, 1951; Hammond, 1998), a sufficient number of rational, independently acting participants (Stigler, 1957), information symmetry among participants (Mas-Colell, 1982; Robinson, 1934), and no market power for small groups (Godal et al., 2005; Mas-Colell, 1982).

The second Welfare Theorem states that a Pareto efficient equilibrium can be achieved via a competitive equilibrium with endowment (Hammond, 1998; Stiglitz, 1991). Governments, for instance, try to enact this by taxing or subsidizing markets, and consequently push for a redistribution of resources or wealth (Hammond, 1998). Such governmental interference, however, constitutes an external intervention (Coase, 1960). Ronald Coase, in turn, argues in his work that even if markets are not complete, they can achieve Pareto efficiency without external intervention through bargaining among market participants (Coase, 1960, 1974, 1981; Farrell, 1987). Such bargaining is possible if the following criteria are met: (1) Property rights are clearly defined, which means fellow market participants give a property rights owner consent to act in a certain way (Demsetz, 1974), including the right to buy and sell said rights (Makowski & Ostroy, 2001), (2) market participants are rational (Ellingsen & Paltseva, 2016), (3) there are no wealth effects (Hoffman & Spitzer, 1982), i.e., an increase in wealth does not lead to changes in spending and savings behavior (Maki & Palumbo, 2001), (4) all participants are perfectly informed, and (5) there are no or low transaction costs (Ellingsen & Paltseva, 2016).

As the scope of requirements attributed to the two Welfare Theorems is non-exhaustive and could be extended as new problems arise (Stigler, 1957), we will prioritize the adherence to the Coase criteria. In human economies, these criteria are often not met: "There is not a complete set of markets; information is imperfect; […] commodities […] are not homogeneous" (Greenwald and Stiglitz, 1986, p. 259) and "transaction costs are ever with us" (Fox, 2007, p. 388). However, it remains to be seen whether machine economies can be constructed to meet Coase's criteria. We summarize these in Table 1 and use them as the basis for defining technical requirements that can guide the identification, selection, and design of efficient machine economies.

Table 1 Coase's criteria for a Pareto efficient economy

Economic criteria
(1) Clearly defined property rights
(2) Rational market participants
(3) No wealth effects
(4) Perfect information
(5) No or low transaction costs

Technical requirements for an efficient machine economy

In this section, we translate Coase's criteria for Pareto efficiency into technical requirements – specifications or features that need to be implemented to achieve efficiency in economies with machine participation. In "Efficiency challenges in machine economies", we then discuss the feasibility and pertinent challenges of different machine economy types with respect to these requirements.

Firstly, Coase stipulates that property rights need to be clearly defined. To do so, it must be verifiable which rights each entity has regarding which resources, including the authorization to sell or transfer these resources. Moreover, property rights must be protected against illicit transfer, since otherwise confidence and reliance would be undermined (Demsetz, 1966). Shleifer (1994) argues that the critical determinants of well-defined property rights are efficient control structures and enforceable contracts. In a human economy, control structures are efficient if property rights owners can take legal recourse to protect their rights. Contracts are enforceable if the legal rights can be upheld through specified mechanisms. From a more technical perspective, these considerations translate to the need for registries that reliably and objectively record and update data on property rights (Kaplow & Shavell, 2002). This data needs to be accessible to all market participants, while operations that modify them need to be restricted, e.g., to the current owner. Moreover, mechanisms need to be implemented that allow to eliminate and invalidate illegitimate transactions and sanction violations (Green, 2002; Kaplow & Shavell, 2002).

The second Coase criterion requires market participants to act rationally. Slovic (2000) describes actors as rational when they make decisions in their own best interest. He further elaborates that rational decision-making is analytical and follows a certain set of logical rules, for instance, probability theory (Slovic, 2000). Monroe and Maher (1995) add that each actor's preferences must be consistent and that actors must have the capability to evaluate the consequences of alternative courses of action. By doing so, they can determine which choice will advance them most toward their defined goal in the sense of the best personal outcome. Machines participating in an efficient machine economy thus need to be designed in a way that enables them to define clear goals and consistently pursue them (Marwala, 2021; Schmidt & Wagner, 2019). To do so, these machines need to gather information relevant to their goal, input it into a decision engine that follows a consistent logic, and analyze input and output relations of different scenarios (Marwala, 2014). Subsequently, they need to be able to rank these scenarios by their impact on the defined goal(s), and then act on the option that leads to the best outcome (Russell et al., 2015).

Following the third Coase criterion, and in line with the condition of consistent preferences, machines must not fall victim to wealth effects (Hoffman & Spitzer, 1982). Wealth effects describe a situation in which an actor's spending behavior and risk profile change in accordance with their wealth (Lettau & Ludvigson, 2004). In human economies, wealth changes can occur either in an unanticipated or an expected manner. Unanticipated wealth changes, for instance, can result from asset price shocks (Paiella and Pistaferri, 2017) and lead to irrational changes in spending behavior. Expected wealth changes, in turn, typically do not lead to substantial shifts in spending behavior, as these changes can be anticipated and planned for. Rational human actors gradually incorporate expectations of changes in their wealth into their consumption behaviors (Paiella & Pistaferri, 2017). Rational machines, however, could avoid them altogether when they follow consistent and stable preferences and evaluate the consequences of their actions objectively.

Although rational machines may not be subject to wealth effects, the fourth Coase criterion nevertheless requires that they have access to perfect information to evaluate the consequences of their actions and make informed decisions. Perfect information allows for observing probability distributions and using them to calculate risk exposures (Awrey, 2012; Winseck, 2002). To confine the broad scope of perfect information, we draw on competitive theory. It demands that all market participants (producers and customers) have perfect information regarding the price, volume, and quality of any good or service in the economy (Awrey, 2012; Stiglitz, 1989). Translated into technical requirements, this means that every participant in an efficient machine economy needs equal, non-discriminatory access to the same set of information that allows transparency on both pricing and the properties of the goods or services in question.

This information should come with low costs for its acquisition (Caplin & Dean, 2015), processing, and sharing. More generally, the fifth Coase criterion requires that transaction costs are kept as low as possible. Besides costs for information access, processing, and sharing, transaction costs also result from the process of discovering counterparties, aligning economic actors on decisions, actions, negotiations, and the terms of an interaction. Additionally, they can include costs for potential re-negotiations, changes, dispute resolution, and assurances that terms and conditions of agreed interactions will be met (Young, 2013). Transaction costs can be categorized into technological and mental transaction costs. The latter refer to costs of searching for and comparing prices, attributes, and preferences (Szabo, 1999). Even though rational machines have stable preferences, they must make comparative decisions and analyze scenarios. While these processes fit the definition of mental transaction costs, they require computational effort. Hence, a machine's mental transaction costs are inherently technological. For this reason, we focus on technological transaction costs in the following paragraphs.

To derive technical requirements from the need for low technical transaction costs, we refer to the work of Ali et al., (2017), Papaefstathiou and Manifavas (2004), Rezaeibagha and Mu (2019), and Szabo (1999). Overall, technological transaction costs are strongly mediated by scalability, latency, and security requirements (Ali et al., 2017; Papaefstathiou & Manifavas, 2004; Rezaeibagha & Mu, 2019). Scalability requirements, for instance, directly influence the costs of modifying property rights. Any such modification comes with costs for providing connectivity, running databases, managing access and permissions, and offering add-on services (Papaefstathiou & Manifavas, 2004). When systems need to support high numbers of transactions and transaction complexity, these costs are typically higher. Machine economies that do not limit the number of participants accordingly require registries that can scale to a high number of participants and transactions (Ali et al., 2017) while keeping the costs per modification low. As some of these modifications need to be persisted, storage costs also need to be low.

Scalability, latency, and security requirements also play an important moderating role for technological transaction costs in the interactions between participants. In machine economies that feature a large number of participants, for instance, it is crucial to consider the number of interactions and involved participants in a transaction, associated communication time, and costs. The same applies to administrative costs, including items such as billing, invoicing, fees, commissions, etc. From a technical perspective, these economies require systems that facilitate the traceability of economic interactions without issuing unnecessary messages, thus avoiding needless communication and processing.

Table 2 summarizes the technical requirements we discussed in this section. It does not purport to be mutually exclusive and comprehensively exhaustive. Some of the listed requirements could also be assigned to multiple criteria or are interdependent with other requirements. They also do not constitute an exhaustive list for meeting the economic criteria defined by Coase. Lastly, we can and do not claim that any of them are strictly necessary.

Efficiency challenges in machine economies

The presented economic criteria and technical requirements are useful for identifying and addressing key efficiency challenges in machine economies. For instance, property rights may often be comparatively easy to define, and machines can be programmed to be rational and immune to wealth effects. Yet, ensuring perfect information and low transaction costs will often be a daunting goal. Moreover, human participants tend to introduce inefficiencies, such as bounded rationality, opportunistic behavior, or contractual violations (Lumineau et al., 2021). To provide a structured overview of these challenges, we distinguish between four types of machine economies (Fig. 1). These types differ in their interactions, i.e., machine-to-human or machine-to-machine, and governance, i.e., enforced by humans or machines (Katzenbach & Ulbricht, 2019). We included a governance dimension as governance mechanisms are common means to mediate economic inefficiencies, such as information asymmetries, externalities, monopolies, and public goods. They "ensur[e] that participants engage in collective and mutually supportive action, that conflict is addressed, and that [system] resources are acquired and utilized efficiently and effectively" (Provan & Kenis, 2008, p. 231). In the following, we use one specific example for each type to highlight how the challenges manifest.

Type 1 – Machine-to-human economies under human governance

Example Energy markets with mixed human–machine trading

A key challenge in the decarbonization of energy systems is balancing demand with green but intermittent power generation from renewable sources like solar and wind. One way of addressing this balancing challenge is through automated peer-to-peer trading systems that allow local producers of green electricity to sell it to local consumers. A field study by Wörner et al. (2022) analyzed such a peer-to-peer trading system. It enabled human participants to specify certain trading parameters which set limits for an auctioning mechanism that automatically balanced local demand and green-energy supply from local producers. In case of imbalances, the local utility company would utilize an algorithm to buy or sell the outstanding difference, at a feed-in or retail price, respectively. In essence, this example describes a machine-to-human economy. As the human programmers behind the auctioning algorithm can change the market/matching mechanisms (provided they stay within the bounds of the law), they directly influence the coordination of the interacting parties and thus exercise human governance.

Table 2 Exemplary technical requirements for a Pareto efficient machine economy

Economic criteria	Exemplary technical requirements
Clearly defined property rights	Registry for property rights and respective owners
	Unrestricted verifiability of property rights and ownership relations
	Restricted modification rights
Rational market participants/No wealth effects	Mechanism to clearly define and identify goals
	Consistent and stable preferences
	Ability to access and filter information
	Capacity to process scenarios via decision engines
	Ability to rank calculated scenarios by impact on defined goals
Perfect information	Non-discriminatory access to a joint set of relevant information for all participants
	Ability to acquire and share information with all users
No or low transaction costs	Scalable registries with low processing costs
	Minimization of storage requirements
	Low complexity of interactions between machines

Peer-to-peer trading systems can be equipped with a clear definition of property rights and their respective owners, such as for each kWh sold and bought by each participant in each trading period. Moreover, the system allows for the definition of consistent goals (for instance, reducing electricity costs, securing energy supply, or maximizing profit). Yet, the human buyers involved in the study showed irrational behavior in multiple ways. They behaved inconsistently regarding their preferences, as they expressed willingness to pay a price premium for green electricity. However, they never actually placed such bids in practice, which suggests indication-behavior gaps (Wörner et al., 2022) and, thus, inconsistency towards their goals. Similarly, human sellers were not willing to offer electricity at prices below the utility company's feed-in tariffs. These various forms of irrational behavior led to many mismatched bids and, therefore, inefficient energy allocation. Moreover, they resulted in negative externalities when the utility company had to step in and sell 'non-green' electricity to balance supply shortages. This signals that neither side of the human participants could compare scenarios adequately, estimate the impact on their goals, and act accordingly.

This example illustrates that efficiency is difficult to achieve in machine economies when human participants

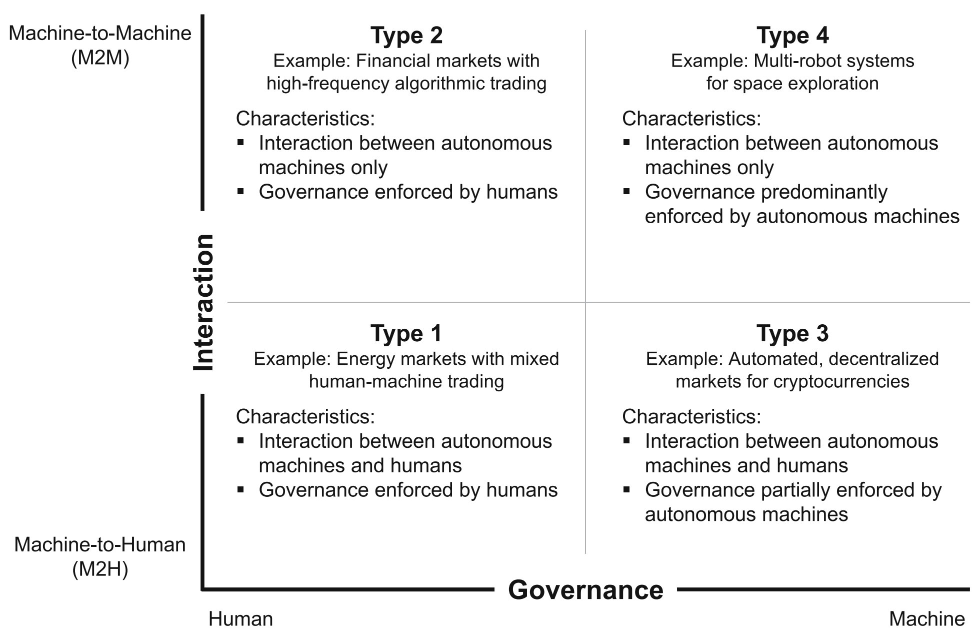

Fig. 1 Machine economy framework

are involved. Mediating irrational human behavior can be challenging and will often require human governance that anticipates or reacts to human 'inefficiencies'.

Type 2 – Machine-to-machine economies under human governance

Example Financial markets with high-frequency algorithmic trading

Algorithmic trading describes a set of systems, such as in-trade-execution programs, that use complex algorithms to automatically analyze, signal, execute, and manage trades in high-liquidity financial markets (Nuti et al., 2011; Treleaven et al., 2013). As the name implies, investors specify their trading goals and strategies in the form of algorithms (or mathematical instructions), which are then carried out by computers. The algorithms can typically optimize order sizes and the timing of trades completely independent of human involvement (Kissell, 2013). In many financial markets, algorithmic trading agents account for enough liquidity and trading volume (Hendershott et al., 2021; Moriyasu et al., 2018) that they often trade among each other, effectively making high-frequency algorithmic trading a machine-to-machine economy. However, humans influence many technical and legal boundary conditions under which algorithmic trading agents operate and thus enforce human governance.

Trading in such markets usually relies on centralized order books (Nuti et al., 2011), which list all buy and sell orders. These order books provide an overview of clearly defined property rights and restrict changes to the respective property rights holders. The trading algorithms, in turn, are typically rational as they are programmed with clearly defined goals and consistent and stable preferences, which also avoids wealth effects. Perfect information, in turn, is a more difficult criterion to achieve. Although technically speaking, trading platforms grant every participant non-discriminatory access to information in the system, there are differences in how fast the algorithms can gather information and derive actions. To increase speed, many investors thus engage in 'arms races' to secure nano- or picosecond leads in information acquisition, e.g., by investing in high-speed fiber-optic cables and reducing the physical distance between their servers and those of the exchanges (Levens, 2015; MacKenzie, 2021; Ye et al., 2013). These investments can lead to significant information asymmetries and eventually mono- or oligopolistic tendencies. Regarding transaction costs, high-liquidity markets are typically "associated with fast trade execution and low transaction costs" (Nuti et al., 2011, p. 67). However, policing and enforcement typically require the involvement of human courts, which are both slow and costly. An illustrative example in this regard is the May 2010 flash crash (Bridegan & Moussa, 2016; Kirilenko et al., 2017), in which the developer of an algorithmic trading agent was only successfully prosecuted years after defrauding and manipulating the market using his trading algorithm. Needless to say, the illegitimate trades were not reversed, and defrauded traders were not refunded.

To sum up, financial markets that feature high-frequency algorithmic trading demonstrate that machine-to-machine economies are possible and can achieve a high degree of economic efficiency. However, the perfect information and low transaction cost criteria can be difficult to fulfill when human actors outside of the machine economies introduce power imbalances, and when costs for human policing and enforcement are high.

Type 3 – Machine-to-human economies under machine governance

Example Automated, decentralized markets for cryptocurrencies

Traditional exchanges often utilize market makers to match buy and sell orders (Rust & Hall, 2003). Many decentralized markets for cryptocurrencies, in contrast, make use of so-called automated market makers. Instead of interacting with another human, traders in these markets directly interact with a liquidity pool managed by a smart contract – an algorithm that enables the automated execution of pre-defined rules – that calculates prices based on a price function and the liquidity of the cryptocurrencies within the automated market maker protocol (Gramlich et al., 2023; Mohan, 2022) before automatically executing the trade. These markets are effectively machine-to-human economies. As governance decisions are pre-defined by means of smart contracts, they also constitute a simple form of algorithmic or machine governance (Lumineau et al., 2021).

By design, decentralized markets for cryptocurrencies rely on blockchain registries with an 'append only' structure (Butijn et al., 2020) that ensure that property rights are clearly defined. Smart contracts keep track of or enforce user rights and obligations (Hartwich et al., 2023; Lin & Liao, 2017) while managing the modification rights of the associated records (Kannengießer et al., 2021). Many markets allow only the trade of one currency pair per transaction, which restricts input parameters (i.e., only quantity) and limits preferences. This setup makes interactions with automated market makers relatively simple, which arguably allows to meet rationality requirements. However, even simple transactions can incur high transaction fees based on the network load and transaction volume of the underlying blockchain (Roughgarden, 2021). Ensuring perfect information is difficult as well. Automated market makers are typically non-discriminatory by design and allow all traders to view pricing information before submitting a trade. However, this information can also be misleading, as parties who validate and add new transactions to the blockchain registry can front-run trades by changing the order of the trades they are validating (Daian et al., 2020), which creates information asymmetries. To mitigate the impact of front-running, higher-value trades can be split into smaller ones that are routed through different exchanges (Kulkarni et al., 2022). Yet, efficiently executing these strategies requires the use of algorithms. This need for algorithmic support exacerbates economic inefficiency as it is typically wealthier market participants that can afford to invest in defensive algorithms against front-running. Moreover, market participants with the financial means to build close ties with validators can avoid being front-run or even front-run others, which creates mono- or oligopolistic tendencies. While mitigation measures have been proposed and implemented, for instance, in the form of proposer-builder separation (Yang et al., 2022), these are not directly exercised on a smart contract level. In fact, many recent examples suggest that smart contracts can only enforce rudimentary governance functions and are heavily dependent on the code's quality and the programmer's choices (Praitheeshan et al., 2019; Thurman, 2021).

In effect, decentralized markets for cryptocurrencies again illustrate how transaction costs and inefficiencies introduced by human actors can constrain machine economies. They are also illustrative examples of the limits of machine governance orchestrated by automated market makers and smart contracts.

Type 4 – Machine-to-machine economies under machine governance

Example Multi-robot systems for space exploration

Type 4 machine economies exist today mostly as conceptual ideas. They may nevertheless find practical applications in the near future, for instance, in the context of space exploration (Leitner, 2009; Yliniemi et al., 2014). NASA and ESA are currently preparing new lunar exploratory missions, such as ARTEMIS and Terrae Novae 2030+ (Smith et al., 2020; Vijendran et al., 2021). These missions bring together many nations and private companies that work on a wide range of robots with heterogeneous capabilities (Borowitz, 2019). In space, these robots will need to cooperate as multi-robot systems (Borowitz, 2019). Multi-robot systems describe "multiple autonomous, interacting [robots] that have common or conflicting goals" (Rizk et al., 2019, p. 3). In such systems, robots coordinate efforts to achieve complex tasks that would be impossible for a single robot to complete (Gautam & Mohan, 2012; Parker, 2012). As coordination with humans on Earth might not always be possible (for instance, because of long signal traveling times), and missions are often time-critical, robots of different nations and companies may need to align with one another and make autonomous decisions on the spot. When they exchange both information and value, multi-robot systems can be seen as machine-to-machine economies. Not only could human governance not be practical under the described circumstances, but there is also no international legal framework for space. Hence, space multi-robot systems require a high degree of machine governance.

Moreover, deploying robots into space comes with a very specific set of challenges. For instance, computational resources in terms of hardware and energy and corresponding software capabilities are highly restricted. In addition, space robots are exposed to a very hostile environment (i.e., cosmic rays) that increases the risk of electronic failures (Bogue, 2012; Futaana et al., 2022). This risk is typically managed through replicated information processing, which multiplies resource consumption (Wensley et al., 1978). In space, it is thus essential to meet the identified technical requirements with particularly low resource needs. Technically speaking, multi-robot systems will highly depend on scalability solutions and compression technologies to control technological transaction costs. Technological transaction costs for the discovery of counterparties and the definition of boundary conditions of economic exchange will also require close attention. Efficient digital platforms may help to keep these costs in check (Hein et al., 2020).

To ensure rational robot behavior in various circumstances, programmers may need to hard-code objectives (Turpin et al., 2014) or equip robots with decision engines capable of producing replicable decisions based on the information at hand (Fung et al., 2021; Janiesch et al., 2021; Linardatos et al., 2020; Loyola-Gonzalez, 2019). Perfect information is, again, a challenge as the national space programs behind the missions have strategic interests and sensitive information to protect (Borowitz, 2019). Consequently, perfect information should be understood contextually (Awrey, 2012; Stiglitz, 1989) and not be equated with 'all robots have access to all data'. Yet, the 'public good' challenge of data remains as data is replicable, non-exclusive, and utilizable by multiple machines at once (Hummel et al., 2021),which makes it challenging to enforce property rights over shared data. Therefore, technologies that enable publicly verifiable statements without disclosing underlying data may be useful in these contexts (Garrido et al., 2022). Moreover, space multi-robot systems will need property rights registries that balance unrestricted verifiability with restricted modification rights (Sedlmeir et al., 2022).

Overall, multi-robot systems could have a high potential of achieving Pareto efficiency when they account for the resource challenges in space. Yet, multi-robot systems will likely require a high degree of machine governance. Unpredictable circumstances will require such machine governance to be dynamic, react to potentially changing circumstances, and adapt governance mechanisms if needed, i.e., through intelligent agent technology.

Conclusion and outlook

With this paper, we aim to establish a fundamental understanding of the economic criteria and technical requirements for *efficient* machine economies. In particular, we discuss Pareto efficiency criteria that can underpin such economies. We then translate these criteria into technical requirements and discuss key efficiency challenges based on a framework of four machine economy types (Fig. 1). This framework is also useful for structuring research opportunities, which we outline in this section (for a summary, see Table 3). Overall, we hope to broaden an increasingly technical discussion and build bridges that will facilitate a more interdisciplinary dialogue on interaction and governance aspects of machine economies.

Interaction

As our examples indicate, machine economies with human interactions can easily fall short in achieving Pareto efficiency due to the bounded rationality, opportunistic behavior, or contract violations of their human participants. It will thus be essential to analyze how humans can collaborate with machines and 'delegate' to machines those decisions and tasks that are particularly prone to irrational behavior or wealth effects. Additionally, further investigation is required to identify how humans can be supported in their decisions and interactions without giving up control over identifying and defining (economic) goals. Teodorescu et al. (2021), for instance, offer a fruitful starting point for this investigation by describing how a machine-human partnership can be designed to enhance human performance. Likewise, Fügener et al. (2021) suggest techniques for personalizing artificial intelligence to support individual goals and motives. Moreover, research on "next generation digital platforms"

Table 3 Avenues for future research on efficient machine economies

Research area	Exemplary research questions
Interaction	
M2H interaction	How can humans delegate tasks that are prone to irrational behavior to machines?
	How can machines support humans while humans still maintain control over identifying and defining economic goals?
	Can digital platforms support a dynamic allocation of tasks between humans and machines?
M2M interaction	Which machine learning methods can support machines in establishing clearly defined goals?
	To what extent do self-assigned goals need to be reproducible and explainable to constitute rationality?
	Which type of information will require which level of availability and verifiability to ensure perfect information in a machine economy?
Governance	
Human governance	How can humans utilize machine learning to identify market designs that support efficient machine economies?
	How can machine-supported accountability measures reduce transaction costs and make human governance more efficient?
	To what extent could machines approximate complete contracts, and how would contracts in a machine economy differ from those in a human economy?
Machine governance	Which governance aspects of human economies are applicable in machine settings, and which may require novel, machine-specific designs?
	To what extent can existing platform designs support hybrid, human–machine settings?
	Under what circumstances can machine governance mechanisms be beneficial to human economies?

(Rai et al, 2019, p. iii) may be instructive in identifying mechanisms for splitting tasks between several machines and humans. It will be especially important to determine how such platform setups can support efficient hybrid settings in which task allocation between humans and machines is dynamic.

Machine-to-machine economies, on the other hand, require further research into decision engines that support rational decision-making and the autonomous definition of clear goals for a broad range of scenarios. The discussion on white box vs. black box machine learning can be instructive in this regard as it provides insights into how machine learning models can self-assign reproducible and explainable goals (Arrieta et al., 2020; Fung et al., 2021; Loyola-Gonzalez, 2019). In particular, it will be important to unpack to what extent self-assigned goals need to be reproducible and explainable to constitute rationality and consistent preferences (Linardatos et al., 2020). Further, it will be crucial to identify which types of machine learning methods (unsupervised, semi-supervised, or reinforcement (Kühl et al., 2022)) should be employed to not only establish a clear and autonomous definition of goals but also to advance the compliance with other economic criteria, such as low transaction costs and information asymmetries. Research on generative adversarial network is instructive in this regard. Jiang et al. (2019), for instance, demonstrate that generative adversarial networks can increase efficacy while reducing computational costs – a key aspect in reducing technological transaction costs. Karras et al. (2020), in turn, show that training generative adversarial network models can produce accurate results from limited data.

More research is also required in the area of scalability, especially for machine-to-machine economies in resource-constrained environments. Regarding data storage and processing, it will be essential to analyze trade-offs between storage compression, the degree of information replicability, and data availability. In-depth analysis will also be required into how information can be identified and classified concerning their necessary level of duplication, verifiability, and accessibility.

Governance

Besides more efficient interactions, further research will also be required into more effective governance for machine economies. Human governance will need to level the playing field for machines by making them independent of their owners' wealth or power status, e.g., by developing new market structures that negate these 'outside' factors. Batch auctions that group individual bids for simultaneous execution could present an interesting starting point for such discussions (Budish et al., 2015). Alternatively, it could be interesting to analyze how machine learning can support the discovery of novel market designs suitable for machine economies. Moreover, research into novel, potentially machine-supported accountability measures might help reduce transaction costs and make human governance more efficient. Also, should efficient machine economies achieve or at least approximate a high degree of perfect information, it would be interesting to analyze if and to which extent machines could support the creation of complete contracts, and how contracts in a machine economy would differ compared to those in a human economy.

Machine governance, in turn, is still in its infancy and can only enforce rudimentary rules (Ferreira, 2021; Levy, 2017). It will thus be important to develop (new) approaches for allocating decision rights and the design of accountability measures. In particular, it will be essential to understand what governance aspects of human economies are applicable in machine settings and which may require novel, machine-specific designs (Rossi et al., 2019). Research is especially required regarding the enforceability of rules and property rights in machine-to-machine interactions. Machine economies will require efficient mechanisms that ensure the cooperation and coordination of transaction partners while providing inclusive access and avoiding abuse of control. Moreover, effective mechanisms to reverse (see, e.g., Wang et al., 2022) and potentially sanction illegitimate transactions in property rights records will be essential. Digital machine identities and the machine-verifiability of established trust structures could provide fruitful starting points for these questions (Sedlmeir et al., 2021).

Lastly, deeper insight into the governance mechanisms of platform ecosystems is needed for both machine-to-human and machine-to-machine settings. Further analyses will be required to determine if received wisdom on platform governance can be applied in machine economy settings. Specifically, research should raise questions about the design of governance mechanisms in human–machine hybrid platforms. Which parameters will determine who (human or machine) will hold which governance rights? Will it be effective to split governance responsibilities or hand them over to machines in pursuit of economic efficiency? Also, research should identify if and under which circumstances machine governance mechanisms may be relevant and beneficial to human economies.

Lastly, it may be necessary to explore socio-technical aspects and spill-over effects to consider "the impact on people, infrastructure, technology, processes, culture and organization[s]" (Sony and Naik 2020, p. 8). As often in IS research, it will thus be essential to take an interdisciplinary perspective to analyze the full potential and implications of an *efficient* machine economy.

Acknowledgements This research was funded in part by the Luxembourg National Research Fund (FNR) and PayPal, PEARL grant reference *13342933/Gilbert Fridgen*, as well as the Luxembourg

National Research Fund (FNR) in the FiReSPArX (grant reference *14783405*) and PABLO (grant reference *16326754*) projects. Moreover, we would like to acknowledge the support of the European Union (EU) within its Horizon 2020 programme, project MDOT (Medical Device Obligations Taskforce), Grant agreement *814654* that made this paper possible. We also thank the editor and the anonymous reviewers for their highly valuable and constructive feedback. For the purpose of open access, the authors have applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission

Declarations

Conflict of interest The authors declare that they have no conflict of interest

References

Ali, S. T., Clarke, D., & McCorry, P. (2017). *The nuts and bolts of micropayments: A survey*. Retrieved February 20, 2022, from arXiv:1710.02964

Amazon Web Services. (2022). *Industrial Internet of Things*. Retrieved October 6, 2022, from https://aws.amazon.com/iot/solutions/industrialiot/

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion, 58*, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

Arrow, K. J. (1951). An extension of the basic theorems of classical welfare economics. *Proceedings of the 2nd Berkeley Symposium on Mathematical Statistics and Probability*, (pp. 507–532). University of California Press

Arrow, K. J. (1966). *Economic equilibrium*. International Encyclopedia of the Social Sciences. Retrieved February 20, 2022, from https://apps.dtic.mil/sti/pdfs/AD0630666.pdf

Arrow, K. J., & Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, 265–290. https://doi.org/10.2307/1907353

Arthur, B. W. (2017). Where is technology taking the economy? (Tech. rep.). McKinsey & Company. https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/where-is-technology-taking-the-economy

Awrey, D. (2012). Complexity, innovation, and the regulation of modern financial markets. *Harvard Business Literature Review, 2*, 235.

Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly, 45* (3), 1433–1450. https://doi.org/10.25300/MISQ/2021/16274

Blaug, M. (2007). The fundamental theorems of modern welfare economics, historically contemplated. *History of Political Economy, 39*(2), 185–207. https://doi.org/10.1215/00182702-2007-001

Bogue, R. (2012). Robots for space exploration. *Industrial Robot: An International Journal*. https://doi.org/10.1108/01439911211227872

Bons, R. W., Versendaal, J., Zavolokina, L., & Shi, W. L. (2020). Potential and limits of blockchain technology for networked businesses. *Electronic Markets, 30*(2), 189–194. https://doi.org/10.1007/s12525-020-00421-8

Borowitz, M. (2019). Strategic implications of the proliferation of space situational awareness technology and information: Lessons learned from the remote sensing sector. *Space Policy, 47*, 18–27. https://doi.org/10.1016/j.spacepol.2018.05.002

Bridegan, T., & Moussa, D. (2016). US v. Sarao: The flash crash and a new effort to prosecute market manipulation and deceptive trading practices. *Georgetown Law Technology Review, 1*, 396.

Brynjolfsson, E., Rock, D., & Syverson, C. (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. University of Chicago Press.

Buckle, M., & Thompson, J. (2020). *The UK financial system: Theory and practice*. Manchester University Press.

Budish, E., Cramton, P., & Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics, 130*(4), 1547–1621. https://doi.org/10.1093/qje/qjv027

Butijn, B.-J., Tamburri, D. A., & van den Heuvel, W.-J. (2020). Blockchains: A systematic multivocal literature review. *ACM Computing Surveys, 53* (3). https://doi.org/10.1145/3369052

Caplin, A., & Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review, 105*(7), 2183–2203. https://doi.org/10.1257/aer.20140117

Coase, R. H. (1960). The problem of social cost. In *Classic Papers in Natural Resource Economics* (pp. 87–137). Springer. https://doi.org/10.1086/674872

Coase, R. H. (1974). The lighthouse in economics. *The Journal of Law and Economics, 17*(2), 357–376. https://doi.org/10.1086/466796

Coase, R. H. (1981). The Coase theorem and the empty core: A comment. *The Journal of Law and Economics, 24*(1), 183–187. https://doi.org/10.1086/466980

Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., & Juels, A. (2020). Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. *IEEE Symposium on Security and Privacy*. https://doi.org/10.1109/sp40000.2020.00040

Davidson, S., De Filippi, P., & Potts, J. (2016). *Economics of blockchain.* https://doi.org/10.2139/ssrn.2744751

Demsetz, H. (1966). Some aspects of property rights. *The Journal of Law and Economics, 9*, 61–70. https://doi.org/10.1086/466619

Demsetz, H. (1974). Toward a theory of property rights. In: *Classic Papersin Natural Resource Economics* (pp. 163–177). Springer. https://doi.org/10.1057/9780230523210_9

Ellingsen, T., & Paltseva, E. (2016). Confining the Coase theorem: contracting, ownership, and free-riding. *The Review of Economic Studies, 83*(2), 547–586. https://doi.org/10.1093/restud/rdw001

Farrell, J. (1987). Information and the Coase theorem. *Journal of Economic Perspectives, 1*(2), 113–129. https://doi.org/10.1257/jep.1.2.113

Ferreira, A. (2021). Regulating smart contracts: Legal revolution or simply evolution? *Telecommunications Policy, 45* (2), 102081. https://doi.org/10.1016/j.telpol.2020.102081

Flood, M. D., et al. (1991). An introduction to complete markets. *Federal Reserve Bank of St. Louis Review, 73*. https://doi.org/10.20955/r.73.32-57

Fox, G. (2007). The real Coase theorems. *Cato Journal, 27*, 373.

Fügener, A., Grahl, J., Gupta, A., & Ketter, W. (2021). Will humans-in-theloop become borgs? Merits clatfalls of working with AI. *MIS Quarterly, 45*. https://doi.org/10.25300/MISQ/2021/16553

Fung, P. L., Zaidan, M. A., Timonen, H., Niemi, J. V., Kousa, A., Kuula, J., Luoma, K., Tarkoma, S., Petäjä, T., Kulmala, M., et al. (2021). Evaluation of white-box versus black-box machine learning models in estimating ambient black carbon concentration. *Journal of Aerosol Science, 152*, 105694. https://doi.org/10.1016/j.jaerosci.2020.105694

Futaana, Y., Shimoyama, M., Wieser, M., Karlsson, S., Andersson, H., Nilsson, H., Wang, X.-D., Fedorov, A., Andre, N., Holmström, M., et al. (2022). Galactic cosmic rays at Mars and Venus: Temporal variations from hours to decades measured as the background signal of onboard microchannel plates. *The Astrophysical Journal, 940*(2), 178. https://doi.org/10.3847/1538-4357/ac9a49

Garrido, G. M., Sedlmeir, J., Uludağ, Ö., Alaoui, I. S., Luckow, A., & Matthes, F. (2022). Revealing the landscape of privacy-enhancing technologies in the context of data markets for the IoT: A systematic literature review. *Journal of Network and Computer Applications, 207*, 103465. https://doi.org/10.1016/j.jnca.2022.103465

Gautam, A., & Mohan, S. (2012). A review of research in multi-robot systems. *IEEE 7th International Conference on Industrial and Information Systems*. https://doi.org/10.1109/ICIInfS.2012.6304778

Godal, O., et al. (2005). Strategic markets in property rights without pricetakers. *Working Paper05/05, Department of Economics, University of Bergen*.

Gramlich, V., Guggenberger, T., Principato, M., Schellinger, B., & Urbach, N. (2023). A multivocal literature review of decentralized finance. *Current knowledge and future research avenues Electronic Markets, 33*, 11. https://doi.org/10.1007/s12525-023-00637-4

Green, S. P. (2002). Plagiarism, norms, and the limits of theft law: Some observations on the use of criminal sanctions in enforcing intellectual property rights. *Hastings LJ, 54*, 167. https://doi.org/10.2139/ssrn.315562

Greenwald, B. C., & Stiglitz, J. E. (1986). Externalities in economies with imperfect information and incomplete markets. *The Quarterly Journal of Economics, 101*(2), 229–264. https://doi.org/10.2307/1891114

Hammond, P. (1998). The efficiency theorems and market failure. *Elements of General Equilibrium Analysis*, 211–260.

Hartwich, E., Ollig, P., Fridgen, G., & Rieger, A. (2023). Probably something: A multi-layer taxonomy of non-fungible tokens. *Internet Research*. https://doi.org/10.1108/INTR-08-2022-0666

Hein, A., Schreieck, M., Riasanow, T., Setzke, D. S., Wiesche, M., Böhm, M., & Krcmar, H. (2020). Digital platform ecosystems. *Electronic Markets, 30*(1), 87–98. https://doi.org/10.1007/s12525-019-00377-4

Hendershott, T., Zhang, X., Zhao, J. L., & Zheng, Z. (2021). Fintech as a game changer: Overview of research frontiers. *Information Systems Research, 32* (1). https://doi.org/10.1287/isre.2021.0997

Hicks, J. R. (1941). The rehabilitation of consumers' surplus. *The Review of Economic Studies, 8*(2), 108–116. https://doi.org/10.2307/2967467

Hoffman, E., & Spitzer, M. L. (1982). The Coase theorem: Some experimental tests. *The Journal of Law and Economics, 25*(1), 73–98. https://doi.org/10.1086/467008

Hummel, P., Braun, M., & Dabrock, P. (2021). Own data? ethical reflections on data ownership. *Philosophy & Technology, 34*(3), 545–572. https://doi.org/10.1007/s13347-020-00404-9

Janiesch, C., Zschech, P., & Heinrich, K. (2021). Machine learning and deep learning. *Electronic Markets, 31*(3), 685–695. https://doi.org/10.1007/s12525-021-00475-2

Jiang, J., Sell, D., Hoyer, S., Hickey, J., Yang, J., & Fan, J. A. (2019). Freeform diffractive metagrating design based on generative adversarial networks. *ACS Nano, 13*(8), 8872–8878. https://doi.org/10.1021/acsnano.9b02371

Jöhnk, J., Albrecht, T., Arnold, L., Guggenberger, T., Lämmermann, L., Schweizer, A., & Urbach, N. (2021). The rise of the machines: Conceptualizing the machine economy. *Proceedings of the 25th Pacific Asia Conference on Information Systems*. https://aisel.aisnet.org/pacis2021/54

Kannengießer, N., Lins, S., Sander, C., Winter, K., Frey, H., & Sunyaev, A. (2021). Challenges and common solutions in smart contract development. *IEEE Transactions on Software Engineering*. https://doi.org/10.1109/tse.2021.3116808

Kaplow, L., & Shavell, S. (2002). Economic analysis of law. In *Handbook of public economics* (pp. 1661–1784). Elsevier. https://doi.org/10.1016/S1573-4420(02)80029-5

Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., & Aila, T. (2020). Training generative adversarial networks with limited data. *Proceedings of Advances in Neural Information Processing Systems, 33*, 12104–12114. https://papers.nips.cc/paper/2020/file/8d30aa96e72440759f74bd2306c1fa3d-Paper.pdf

Katzenbach, C., & Ulbricht, L. (2019). Algorithmic governance. *Internet Policy Review, 8* (4), 1–18. https://doi.org/10.14763/2019.4.1424

Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The flash crash: High-frequency trading in an electronic market. *The Journal of Finance, 72*(3), 967–998. https://doi.org/10.1111/jofi.12498

Kissell, R. (2013). *The science of algorithmic trading and portfolio management*. Academic Press.

Kühl, N., Schemmer, M., Goutier, M., & Satzger, G. (2022). Artificial intelligence and machine learning. *Electronic Markets*. https://doi.org/10.1007/s12525-022-00598-0

Kulkarni, K., Diamandis, T., & Chitra, T. (2022). Towards a theory of maximal extractable value I: Constant function market makers. Retrieved January 31, 2023, from arXiv:2207.11835

Lee, S. M., Olson, D. L., & Trimi, S. (2010). Strategic innovation in the convergence era. *International Journal of Management and Enterprise Development, 9* (1). https://doi.org/10.1504/IJMED.2010.035304

Leiding, B., Sharma, P., & Norta, A. (2021). The machine-to-everything (M2X) economy: Business enactments, collaborations, and e-governance. *Future Internet, 13*(12), 319. https://doi.org/10.3390/fi13120319

Leitner, J. (2009). Multi-robot cooperation in space: A survey. *Advanced Technologies for Enhanced Quality of Life, 144–151*,. https://doi.org/10.1109/AT-EQUAL.2009.37

Lettau, M., & Ludvigson, S. C. (2004). Understanding trend and cycle in asset values: Reevaluating the wealth effect on consumption. *American Economic Review, 94* (1), 276–299. https://doi.org/10.1257/000282804322970805

Levens, T. E. (2015). Too fast, too frequent? High-frequency trading and securities class actions. *The University of Chicago Law Review, 82* (3), 1511–1557. https://www.jstor.org/stable/43575203

Levy, K. E. (2017). Book-smart, not street-smart: Blockchain-based smart contracts and the social workings of law. *Engaging Science, Technology, and Society, 3*. https://doi.org/10.17351/ests2017.107

Lin, I.-C.,& Liao, T.-C. (2017). A survey of blockchain security issues and challenges. *International Journal of Network Security, 19*(5), 653–659. https://doi.org/10.6633/IJNS.201709.19(5).01

Linardatos, P., Papastefanopoulos, V., & Kotsiantis, S. (2020). Explainable AI: A review of machine learning interpretability methods. *Entropy, 23*(1), 18. https://doi.org/10.3390/e23010018

Loyola-Gonzalez, O. (2019). Black-box vs. white-box: Understanding their advantages and weaknesses from a practical point of view. *IEEE Access, 7*, 154096–154113. https://doi.org/10.1109/ACCESS.2019.2949286

Lumineau, F., Wang, W., & Schilke, O. (2021). Blockchain governance–a new way of organizing collaborations? *Organization Science*, 32 (2), 500 – 521. https://doi.org/10.1287/orsc.2020.1379

MacKenzie, D. (2021). Trading at the speed of light. In: *Trading at the speed of light*. Princeton University Press. https://doi.org/10.1515/9780691217796

Maki, D. M., & Palumbo, M. G. (2001). Disentangling the wealth effect: A cohort analysis of household saving in the 1990s. https://doi.org/10.2139/ssrn.268957

Makowski, L., & Ostroy, J. M. (2001). Perfect competition and the creativity of the market. *Journal of Economic Literature, 39*(2), 479–535. https://doi.org/10.1257/jel.39.2.479

Marwala, T. (2014). *Artificial intelligence techniques for rational decision making*. Springer.

Marwala, T. (2021). *Rational machines and artificial intelligence*. Academic Press.

Mas-Colell, A. (1982). *Noncooperative approaches to the theory of perfect competition: Presentation*. In Noncooperative approaches to the theory of perfect competition: Elsevier. https://doi.org/10.1016/B978-0-12-476750-8.50005-1

Mazzucato, M. (2016). From market fixing to market-creating: A new framework for innovation policy. *Industry and Innovation, 23*(2), 140–156. https://doi.org/10.1080/13662716.2016.1146124

Mohan, V. (2022). Automated market makers and decentralized exchanges: A DeFi primer. *Financial Innovation, 8* (1). https://doi.org/10.1186/s40854-021-00314-5

Monroe, K. R., & Maher, K. H. (1995). Psychology and rational actor theory. *Political psychology, 16* (1). https://doi.org/10.2307/3791447

Moriyasu, H., Wee, M., & Yu, J. (2018). The role of algorithmic trading in stock liquidity and commonality in electronic limit order markets. *Pacific-Basin Finance Journal, 49*, 103–128. https://doi.org/10.1016/j.pacfin.2018.04.004

Nuti, G., Mirghaemi, M., Treleaven, P., & Yingsaeree, C. (2011). Algorithmic trading. *Computer, 44*(11), 61–69. https://doi.org/10.1109/MC.2011.31

Paiella, M., & Pistaferri, L. (2017). Decomposing the wealth effect on consumption. *Review of Economics and Statistics, 99*(4), 710–721. https://doi.org/10.1162/REST_a_00629

Papaefstathiou, I., & Manifavas, C. (2004). Evaluation of micropayment transaction costs. *Journal of Electronic Commerce Research, 5*(2), 99–113. http://www.jecr.org/sites/default/files/05_2_p03.pdf

Parker, L. E. (2012). Decision making as optimization in multi-robot teams. *International Conference on Distributed Computing and Internet Technology*, 35–49. https://doi.org/10.1007/978-3-642-28073-3_4

Praitheeshan, P., Pan, L., Yu, J., Liu, J., & Doss, R. (2019). Security analysis methods on ethereum smart contract vulnerabilities: A survey. https://doi.org/10.48550/arXiv:1908.08605

Provan, K. G., & Kenis, P. (2008). Modes of network governance: Structure, management, and effectiveness. *Journal of Public Administration Research and Theory, 18*(2), 229–252. https://doi.org/10.1093/jopart/mum015

Rai, A., Constantinides, P., & Sarker, S. (2019). Next generation digital platforms: Toward human-AI hybrids. *MIS Quarterly, 43*(1), iii–ix.

Rezaeibagha, F., & Mu, Y. (2019). Efficient micropayment of cryptocurrency from blockchains. *The Computer Journal, 62*(4), 507–517. https://doi.org/10.1093/comjnl/bxy105

Rizk, Y., Awad, M., & Tunstel, E. W. (2019). Cooperative heterogeneous multi-robot systems: A survey. *ACM Computing Surveys, 52* (2). https://doi.org/10.1145/3303848

Robinson, J. (1934). What is perfect competition? *The Quarterly Journal of Economics, 49*(1), 104–120. https://doi.org/10.2307/1883878

Rossi, M., Mueller-Bloch, C., Thatcher, J. B., & Beck, R. (2019). Blockchain research in information systems: Current trends and an inclusive future research agenda. *Journal of the Association for Information Systems*,1388 – 1403. https://doi.org/10.17705/1jais.00571

Roughgarden, T. (2021). Transaction fee mechanism design. *ACM SIGecom Exchanges, 19*(1), 52–55. https://doi.org/10.1145/3476436.3476445

Russell, S., Dewey, D., & Tegmark, M. (2015). Research priorities for robust and beneficial artificial intelligence. *AI Magazine, 36*(4), 105–114. https://doi.org/10.1609/aimag.v36i4.2577

Rust, J., & Hall, G. (2003). Middlemen versus market makers: A theory of competitive exchange. *Journal of Political Economy, 111*(2), 353–403. https://doi.org/10.1086/367684

Schmidt, C. G., & Wagner, S. M. (2019). Blockchain and supply chain relations: A transaction cost theory perspective. *Journal of Purchasing and Supply Management, 25*(4), 100552. https://doi.org/10.1016/j.pursup.2019.100552

Schreieck, M., Wiesche, M., & Krcmar, H. (2016). Design and governance of platform ecosystems- key concepts and issues for future research. *Proceedings of the 24th European Conference on Information Systems*. https://aisel.aisnet.org/ecis2016_rp/76

Sedlmeir, J., Lautenschlager, J., Fridgen, G., & Urbach, N. (2022). The transparency challenge of blockchain in organizations. *Electronic Markets, 32*, 1779–1794. https://doi.org/10.1007/s12525-022-00536-0

Sedlmeir, J., Smethurst, R., Rieger, A., & Fridgen, G. (2021). Digital identities and verifiable credentials. *Business & Information Systems Engineering, 63*(5), 603–613. https://doi.org/10.1007/s12599-021-00722-y

Shleifer, A. (1994). Establishing property rights. *The World Bank Economic Review, 8*, 93–117. https://doi.org/10.1093/wber/8.suppl_1.93

Slovic, P. (2000). Rational actors and rational fools: The influence of affect on judgment and decision-making. *Roger Williams UL Review, 6*, 163.

Smith, M., Craig, D., Herrmann, N., Mahoney, E., Krezel, J., McIntyre, N., & Goodliff, K. (2020). The Artemis program: An overview of NASA's activities to return humans to the Moon. *IEEE Aerospace Conference*. https://doi.org/10.1109/AERO47225.2020.9172323

Sony, M., & Naik, S. (2020). Industry 4.0 integration with socio-technical systems theory: A systematic review and proposed theoretical model. *Technology in Society, 61*, 101248. https://doi.org/10.1016/j.techsoc.2020.101248

Stigler, G. J. (1957). Perfect competition, historically contemplated. *Journal of Political Economy, 65* (1). https://doi.org/10.1086/257878

Stiglitz, J. E. (1981). The allocation role of the stock market: Pareto optimality and competition. *The Journal of Finance, 36*(2), 235–251.

Stiglitz, J. E. (1989). Imperfect information in the product market. *Handbook of Industrial Organization, 1*, 769–847. https://doi.org/10.1016/S1573-448X(89)01016-2

Stiglitz, J. E. (1991). *The invisible hand and modern welfare economics.* https://doi.org/10.3386/w3641

Szabo, N. (1999). Micropayments and mental transaction costs. *2nd Berlin Internet Economics Workshop, 44*, 44.

Teodorescu, M. H., Morse, L., Awwad, Y., & Kane, G. C. (2021). Failures of fairness in automation require a deeper understanding of human-ml augmentation. *MIS Quarterly, 45* (3). https://doi.org/10.25300/MISQ/2021/16535

Thurman, A. (2021). *Compound founder says $80M bug presents ' moral dilemma' for DeFi users*. https://www.coindesk.com/tech/2021/10/01/compound-founder-says-80m-bug-presents-moral-dilemma-for-defi-users/

Treleaven, P., Galas, M., & Lalchand, V. (2013). Algorithmic trading review. *Communications of the ACM, 56*(11), 76–85. https://doi.org/10.1145/2500117

Turpin, M., Mohta, K., Michael, N., & Kumar, V. (2014). Goal assignment and trajectory planning for large teams of interchangeable robots. *Autonomous Robots, 37*(4), 401–415. https://doi.org/10.1007/s10514-014-9412-1

Vijendran, S., Schlutz, J., Gerst, A., Istasse, E., De Mey, S., & Schmitt, D. (2021). ESA Terrae Novae exploration strategy 2040. *European Planetary Science Congress*, EPSC2021–798.

Wang, K., Wang, Q., & Boneh, D. (2022). ERC-20R and ERC-721R: Reversible transactions on Ethereum. arXiv:2208.00543

Wensley, J. H., Lamport, L., Goldberg, J., Green, M. W., Levitt, K. N., Melliar-Smith, P. M., Shostak, R. E., & Weinstock, C. B. (1978). SIFT: Design and analysis of a fault-tolerant computer for aircraft control. *Proceedings of the IEEE, 66*(10), 1240–1255. https://doi.org/10.1109/PROC.1978.11114

Winseck, D. (2002). Illusions of perfect information and fantasies of control in the information society. In *Citizenship and Participation in the Information Age* (pp. 33–55). University of Toronto Press. https://doi.org/10.1177/14614440222226280

Wörner, A., Tiefenbeck, V., Wortmann, F., Meeuw, A., Ableitner, L., Fleisch, E., & Azevedo, I. (2022). Bidding on a peer-to-peer energy market: An exploratory field study. *Information Systems Research*. https://doi.org/10.1287/isre.2021.1098

Yang, S., Zhang, F., Huang, K., Chen, X., Yang, Y., & Zhu, F. (2022). SoK: MEV countermeasures: Theory and practice. arXiv:2212.05111

Ye, M., Yao, C., & Gai, J. (2013). The externalities of high frequency trading. *WBS Finance Group Research Paper*, (180). https://doi.org/10.2139/ssrn.2066839

Yliniemi, L., Agogino, A. K., & Tumer, K. (2014). Multirobot coordination for space exploration. *AI Magazine, 35*(4), 61–74. https://doi.org/10.1609/aimag.v35i4.2556

Young, S. D. (2013). Transaction cost economics. In S. O. Idowu, N. Capaldi, L. Zu, & A. D. Gupta (Eds.), *Encyclopedia of Corporate Social Responsibility* (pp. 2547–2552). https://doi.org/10.1007/978-3-642-28036-8_221