

Finance, Technology and Disruption

Jiafu An^{a,b} and Raghavendra Rau^c

^a *University of Edinburgh Business School, University of Edinburgh, United Kingdom*

^b *Department of Accounting and Financial Management, University of Portsmouth, Portsmouth,
United Kingdom*

^c *Judge Business School, University of Cambridge, United Kingdom*

Abstract

In this paper, we assess how recent technology advances have changed the way we coordinate. After a brief discussion of the common challenges to effective coordination, we highlight some important implications of technology on addressing informational and behavioral frictions. We focus on discussing the effects of three specific technology developments including artificial intelligence (AI), automation, and blockchain, on the choice of coordination modes. We argue that technology is shifting the boundaries between firms and markets and is opening the door to new research directions.

Keywords: Coordination, Technology, Boundaries of firm, AI, Automation, Blockchain

1. Introduction

Over the past several decades, academic research has consistently emphasized the role of coordination in economic transactions. Examples include Commons (1931), who stresses that all types of engagement between human beings are based on coordinated transactions, and Hayek (1933) who argues that the operation of economic systems occurs through market coordination. The two most common solutions to the coordination problem have been the use of markets and firms.¹ Many economists have sought to understand where the boundaries between firms and markets lie (Coase 1937; Arrow 1969; Williamson 1971, 1973; Alchian and Demsetz 1972). In his seminal work on the nature of the firm, for example, Coase (1937) sets up a theoretical framework to show that a task should be coordinated within firms when transaction costs on markets outweigh that in firms. Extant research argues that boundaries between firms and markets are dynamic (Coase, 1960; Williamson 1971), since many factors, such as technology advances, can disproportionately change the transaction costs within firms and markets and break the current equilibria in coordination modes.

In this paper, we assess how recent technology advances, including automation, artificial intelligence (AI), and blockchain, have changed the way we coordinate. In other words, we ask how technological progress has been moving the boundaries between firms and markets over the past decade. Given the striking advances in technology during the past decades and the disruptions they have caused in many industries (Brynjolfsson and McAfee 2014; Ford 2015), it is important to examine their impact on forms of coordination, which ultimately dictate the allocation of resources, the process of production, and transactions in our society.

This paper differs from, but also complements, a number of recent studies that examine the impact of technological developments on various economic outcomes, such as labor demand and productivity (Acemoglu and Restrepo 2018, 2019a; Graetz and Michaels 2018), wage and inequality (Autor, Levy and Murnane 2003; Goos and Manning 2007; Michaels, Natraj and Van Reenen 2014), entrepreneurial finance (Howell, Niessner and Yermack 2018; Yermack 2018), corporate governance (Yermack 2017), central banking (Raskin and Yermack 2016) and judicial decisions (Kleinberg et al. 2017, 2018), among others. Distinct from these

¹ We limit our discussion to these two common forms of coordination and aim to provide general implications of the change in technology on coordinating economic activities. Other important coordination mechanisms, including vertical integration (Williamson, 1971) and coordination through various forms of committees (Farrell and Saloner, 1988), among others, are outside of the scope of this discussion.

studies that concentrate on specific economic activities, our paper focuses on the fundamental coordination problem underlying all economic behaviors (Commons 1931). For example, the finding in Graetz and Michaels (2018) that automation using robots increases annual labor productivity growth is at least partly driven by a reduction in coordination costs within firms. Similarly, the emergence of initial coin offerings (ICOs) as an entrepreneurial financing instrument can be understood as a shift of the coordination task from a limited pool of investors (e.g., banks, hedge funds, and other institutions) to literally all potential investors with an internet access (Howell, Niessner and Yermack 2018). Even judicial decisions on bail and discrimination, traditionally coordinated by a specific judge, can now be coordinated by a much larger collective wisdom at little cost via machine learning algorithms (Kleinberg et al. 2017, 2018). Our paper complements these studies by discussing how technology changes the way we coordinate and how this impacts various economic outcomes.

The rest of this paper is organized as follows. Section 2 describes the two coordination approaches, namely coordination through a market or through a hierarchy (firm), discusses the common challenges of effective coordination, and presents some general conditions under which one approach is preferred to the other. Section 3 discusses how tools created by recent technology advances facilitate effective coordination and discusses the implications of specific technological developments, including automation, artificial intelligence (AI), and blockchain, on the way we coordinate. Section 4 concludes.

2. Two common approaches to coordination: Markets versus hierarchies

In this section, we begin by introducing the two most common approaches to coordination, namely coordination through markets and hierarchies (firms). We then discuss the common challenges in effective coordination followed by the general conditions under which one approach supersedes the other.

2.1 Concepts of market and hierarchy coordination

Coordination through the markets is often characterized by economists as a voluntary exchange where price serves to coordinate activities, allocate resources, and provide incentives. Sir Arthur Salter (Plant 1932) vividly describes the economic system as follows:

“The normal economic system works itself. For its current operation, it is under no central control, it needs no central survey. Over the whole range of human activity and human

need, supply is adjusted to demand, and production to consumption, by a process that is automatic, elastic and responsive.”

The voluntary, automatic nature of market coordination has been repeatedly emphasized by prominent economists such as Robertson (1923), Robbins (1932) and Hayek (1933), among others. However, as noted in Coase (1937), this does not necessarily mean that individuals do not choose how to coordinate. In a competitive environment with transaction costs, individuals actively choose between alternative approaches of coordination. One such important alternative is coordination via hierarchies or firms².

Within a hierarchy, activity is coordinated and resources are allocated at least partly by command. Market transactions regulated by price movements are reduced in importance³ and substituted by an entrepreneur-coordinator within a firm, who allocates resources and directs the production process (Coase 1937). Decision makers, i.e., managers and department heads, have broad legal rights to determine how resources are used within the organization.

2.2 Common challenges of effective coordination

Extant research has highlighted several cost advantages in coordination through hierarchies (Coase 1937; Usher 1920; Carlton 1976). For example, the costs of discovering relevant prices of various production factors, or negotiating and concluding a separate contract for each transaction are typically lower within a firm than in a market.

Nonetheless, coordinating through firms also faces several challenges. For instance, as the size of the firm increases, the costs of organizing additional transactions and allocating resources efficiently will rise. This is the so-called diminishing return to management problem (Klador 1934; Robinson 1934). As illustrated in Coase’s (1937) framework, there are decreasing returns to management when the number of transactions organised in a firm rises. For example, entrepreneurs may fail to allocate factors of production to places where they can generate the greatest value as firm size increases. Moreover, as a firm becomes larger, the costs of organising additional transactions within the firm may also rise. Therefore, the choice of whether to coordinate through markets or firms depends critically on the relationship between the marginal transaction costs in the firm and the costs of carrying out the same transaction in the open market.

² Henceforth, we use coordination through hierarchies and firms interchangeably.

³ Even within a hierarchy, not all resource allocations are coordinated through command. For example, in Lancashire cotton industry, looms and yarn can be obtained on credit (Coase 1937).

In addition, the problems of assignment (who gets to decide and when), incentives (how should individuals be compensated) and performance evaluation (how should performance of individuals within an organisation be evaluated) take on prominence, if activity is coordinated within a firm.

Williamson (1973) stresses that the determinants of transaction costs are largely the same in firms and in markets. In his words, “substantially, the same factors that are ultimately responsible for market failures also explain failures of internal organization” (Williamson 1973, p.316). He then summarizes these determinants as the frequency with which transactions are carried out, the specificity of the asset committed to the transaction, the extent of uncertainty over which it becomes difficult to coordinate, the types of opportunistic behavior that people are tempted to commit both before and during contract execution and renewal, and the limited capacities of individuals to receive, store, retrieve and process information without error. The last factor, limited capacity, is also known as bounded rationality. Miller (1956) emphasizes that the human brain is not good at processing huge amounts of information. Humans are only able to juggle about a half a dozen distinct pieces of information at the same time. Williamson (1973) argues that transactions should be coordinated within firms when these factors are prominent.

To facilitate our discussion of the impacts of technological advances on coordination approaches, we further categorize these determinants of transaction costs into three common challenges of effective coordination, namely imperfect information, asymmetric information and behavioral biases. Roughly speaking, imperfect information refers to insufficient information; asymmetric information corresponds one side of a contract having more information than the other; and behavioral biases refer to situations where we are unable to process the information we have in an efficient way.

Imperfect information can incorporate, but is not limited to the issues of uncertainty, specificity, and opportunistic behaviour discussed above. For example, if the specificity of the assets committed to a transaction is high, then the potential for opportunistic behavior due to a lack of information is likely to be mitigated. Similarly, asymmetric information can also include factors such as specificity and opportunistic behavior. Behavioral biases can also incorporate factors such as bounded rationality, frequency, and specificity.

2.3 Boundaries between firms and markets

The main reason that one coordination approach is superseded by the other in some situations is the relative level of transaction costs involved in these approaches (Coase 1937; Williamson 1973). Although both approaches incur costs of imperfect information, asymmetric information and behavioral biases, the relative weights of these costs are different across the two coordination mechanisms. In particular, while information problems are relatively more severe in the price mechanism⁴, behavioral biases, such as bounded rationality, are likely to be comparatively more important in hierarchy coordination.

3. Technology advances and the changing boundaries between firms and markets

In this section, we first outline potential coordination problems in the age of big data after a short review of the common challenges for effective coordination. We then discuss how technology advances, including improved data ontology (i.e., the way that data is labelled), advanced matching algorithms, and preference-identification machine learning systems, can help us mitigate these frictions. Lastly, we describe the implications of specific technology advances, including automation, AI, and blockchain on the choices of coordination methods.

The three common problems of coordination, as discussed in section 2.2 in detail, include imperfect information, asymmetric information and behavioral biases. The important question we ask here is how improvements in technology are changing the way we approach these problems?

3.1 The problems of big data

The recent expansion of large observational datasets on social and economic activity with near-universal coverage makes data ontology, which refers to the way of effectively labelling data⁵, increasingly important (Einav and Levin 2014). On the one hand, the expansion of quality and quantity of data offers market participants with more information, and the increased information flow, in turn, may level up the playing field on both sides of a contract by

⁴ Henceforth, we follow previous literature and use coordination through markets and the price mechanism interchangeably (Coase 1937; Williamson 1973).

⁵ The artificial intelligence literature contains many definitions of data ontology; many of these are contradictory (Noy and McGuinness, 2001). For the purposes of this paper, we define data ontology as a formal explicit description of the characteristics (features and attributes) of observations in a dataset. We refer efficiency of data ontology here as the speed and accuracy of describing a feature or attribute of one or many observations in a given dataset, and the extent to which the end users can identify the observation with the corresponding data ontology at ease.

mitigating the incomplete and asymmetric information problems. This is probably what regulators prefer. For example, in the US, people selling their cars are required to inform buyers of any major accidents the car has been involved in; companies listed on the stock market are required to file quarterly financial reports with the stock market regulator, which are then made public; banks and investment funds must comply with stringent reporting obligations; firms operating in certain sectors (e.g., finance, pharmaceuticals, health care, education, air travel, etc.) are required to provide additional information to regulators and the public.

However, the expansive amount of information also increases the probability of behavioral biases. One important type of resultant behavioral bias is that market participants can unknowingly change their behaviours while keeping their preferences constant, rendering effective policy-making more difficult. For example, consider the “bottomless bowls” experiment in Wansink, Painter and North (2012), where one of two groups of participants with similar body mass index (BMI) are given self-refilling bowls of soup. Despite consuming 73% more soup, the group of participants with self-refiling bowls did not believe that they had consumed more, nor did they perceive themselves as more sated compared with the other group with normal bowls of soup. As economists and regulators often make policy based on observational data, more information can create frictions between the inferred and the actual preferences of market participants.

Another problem created by the expansion of data is information overload. In other words, it provides us too many options that the human brain is not capable of processing efficiently (Miller 1956). For example, suppose you prefer bananas to apples, organic to regular, and ripe to green. How do you choose between green conventional bananas and ripe organic apples? What dimension is more important to you? What happens if there are more dimensions, such as how and where it was grown?

One possible solution to information overload is to condense available information into a single price. But condensing itself has the potential to create additional behavioral biases. One such behavioral bias is price anchoring, where retailers would initially propose a fictional higher price and then mark it down so consumers feel good about the real price (Wansink, Kent, and Hoch 1998). For example, when the iPad was first introduced to the market, Steve Jobs first created an artificial reference point price of \$999 on the iPad after which the real price of \$499 had the effect of making the iPad look inexpensive; the Wall Street Burger Shoppe has burgers priced at \$175, so other items on the menu, including the \$90 steak, would look cheap.

The discussion so far leads us to conclude that, in order to achieve effective coordination, we need (i) a standard language to compare preferences; (ii) a tool to better match preferences along multiple dimensions and (iii) an effective way to comprehensively capture our preferences.

3.2 Solutions brought by technological development

Recent technological developments, including improved data ontology, advanced matching algorithms, and preference-identification machine learnings systems, can help us mitigate the frictions associated with incomplete and asymmetric information as well as behavioral biases. In particular, improvements in data ontology, which refers to a way of effectively classifying and labelling data, can serve as a standard language that allows us to compare preferences across multiple dimensions. One of the problems with the exponential growth in data availability is that much of the information, such as emails, web pages, images, audio and video files, does not fit neatly into a database field. For example, consider YouTube. Uploaders (i.e., sellers) transact with viewers (i.e., buyers) and are often financed by a third group of market participants, advertisers. Viewers need to be able to find desired content easily and content providers need to be able to make their content quickly discoverable. One solution to this problem is to add labels and keywords to the video, but this works only if the uploader picks the right keywords. It is relatively easy to label the features of electronics on a specialized platform, but it is generally very difficult to do so on a general market like YouTube. With improvements in data ontology technology, however, firms are attempting to categorize product information automatically. Start-ups such as Alation, Corrigan, and Expertmaker, among others, are at the frontline of this technology.

In addition, developments in matching algorithms enable us to find and select the optimal transaction partner. As discussed in section 3.1, one feature of the increasingly rich information flow is the difficulty of processing it. In other words, we have too many options to filter and select, and thus to identify the optimal match. However, with advances in technology, matching algorithms are now available to evaluate sets of multiple preferences and their relative weights, and therefore are identifying the best matches for us. The underlying logic is straightforward. Preference data is just another data stream forming a particular pattern. Similar to analysing patterns in regular data, we can adapt pattern-matching algorithms to help us identify optimal transaction partners (Rostek and Yoder 2018). Recent applications of such matching algorithms include, but not limited to, Netflix movie recommendations, Spotify, Apple Music, and Amazon's product recommendations.

Last but not least, developments in machine learning systems can help us in identifying our preferences accurately so that we don't have to spend time making these preferences and their relative weightings explicit. For reasons of information overload and various behavioral biases, it is difficult for us to accurately assess our own preferences, let alone make informed decisions based on the preferences. Fortunately, with huge volumes of data and frequent feedback for the system, machine learning systems can now identify our preferences accurately without the involvement of human social-cognitive skills (Youyou, Kosinski and Stillwell 2015). For example, using a novel dataset that consists 66,732 individuals' written languages on Facebook, Park et al. (2015) build a model that can predict individuals' personalities. They further compare the predicted personality traits with the results of self-reported Big Five personality traits⁶ in a separate sample of 4,824 Facebook users and find that the machine learning-based predictions converge with the self-reported personalities and the validity of the predictions are stable even in a retest after six months. Another example is from a smartphone lending platform, Tala, which offers instant credits (average size of \$50) to people in underrepresented markets such as Kenya, Tanzania, India, and the Philippines. Tala is able to achieve a 90% repayment rate based on a series of behavioral criteria. For instance, if a borrower communicates with more than 58 different contacts, he/she is a good borrower because he/she has a wider network to depend on; if more than 40% of the entries in a borrower's contact list have both first and last names, it suggests that this borrower is 16 times more reliable than one with very few contacts listed with first and last names. These examples suggest that (i) human behavior is predictable⁷ (in aggregated datasets); and (ii) machine learning algorithms are able to facilitate us identifying our preferences accurately.

In summary, improvements in data ontology should help us extract valuable information from huge streams of data and categorize it in many dimensions; advances in matching algorithms should enable us to find and select the optimal transaction partner in the market of our choice; and machine learning systems should be able to identify our preferences as they observe us so that we don't have to spend time making these preferences (and their relative weightings) explicit.

⁶ For more information on the five factor personality traits, see Costa and McCrae (1992).

⁷ Evidence so far shows that human behavior is predictable in isolation. However, much of human behavior also depends on the immediate environment, causing the formation of emergent behaviors. Evidence on whether emergent behaviors are predictable is currently limited.

3.3 Implications of AI, automation and blockchain

In this section, we discuss the implications of three specific technological advances, namely AI, automation, and blockchain, on the way we coordinate. Specifically, we highlight the mechanisms through which each of these technological developments helps us mitigate the problems of effective coordination, namely asymmetric information, incomplete information, and behavioral biases. The channels being emphasized here are discussed in detail in section 3.2, which include data ontology, matching algorithms and machine learning systems.

3.3.1 AI

The starting point of AI is the availability of a large dataset along with a set of labels that describe the features in the data. The goal is to use this data to train a computer so that it can predict labels for a new set of data. For example, one of the AI applications is to train computers to identify images. The conventional approach requires substantial domain expertise and careful engineering to design a feature identifier that transforms raw images into a useful dataset, such as assigning a meaningful and consistent value system to pixels of raw images. Transformed data, which is a suitable internal representation of the raw images, are then used by the learning subsystem to detect or classify patterns in the raw images (LeCun, Bengio and Hinton, 2015). However, this approach was limited in their ability to process natural data in their raw form and achieved limited success. The modern approach works directly with the raw pixels and allows a learning machine to automatically transform raw inputs into meaningful datasets (i.e., images with a suitable internal representation). This has been remarkably successful, not only with image recognition but also with voice recognition, language translation and other traditionally difficult machine learning tasks (Varian, 2014; 2018).

How does AI work? Consider deep learning, a subset of machine learning algorithms. The basic deep learning algorithm allows computational models that are composed of multiple process layers to learn representations of data with multiple levels of abstraction (LeCun, Bengio and Hinton, 2015). Representations of data here refer to the transformation of raw data into an internal suitable, recognizable system of vectors. As emphasize in Witten et al. (2016), with multiple layers of such data transformation, very complex functions can be learned. For classification and identification jobs, higher layers of representation emphasize aspects of the input that are important for discrimination (e.g., between a white wolf and a Samoyed) and suppress irrelevant variations (e.g., different angles of a white wolf). For example, neural networks, which are a form of deep learning systems, consist of several layers, with each layer

trying to discern more and more complex forms in the data. The last layer will assign a probability function to the recognised shape. Figure 1 provides a detailed illustration of how neural networks recognize a dog in a photo. The deep learning algorithms only exploded recently because of the exponential amount of data that can be generated and stored, and the plunge in computing power cost (Nordhaus 2007).

As discussed in previous sections, machine learning algorithms can help us mitigate behavioral biases such as information overload and bounded rationality. In learning our preferences through various data (e.g., consumption) generated by ourselves, we don't have to spend time making these preferences and their relative weightings explicit. As stressed by Youyou, Kosinski and Stillwell (2015), with huge volumes of data and frequent feedback for the system, machine learning algorithms can now identify our preferences accurately without the involvement of human social-cognitive skills.

However deep learning algorithms also have problems. For example, in recurrent neural networks, another type of deep learning algorithm, which allows information to persist, if we change one object in the dataset, the predicted probabilities on everything else will change since they are inter-dependent. This can be better illustrated in an example in Figure 2. The difference in panel A and B is the presence of an elephant in the room (panel B). However, the probabilities of everything else assigned by the same algorithms change dramatically. This problem has significant implications. Consider an autonomous car that goes blind to a pedestrian just because a second earlier it passed a turkey on the side of the road. Other problematic applications, such as Tay, a Microsoft chat bot that repeated racist tweets, and Google's photo tagging algorithms that identified black people as gorillas, are also worth noting.

In sum, with its strong ability to identify patterns in complex data, AI development can help us with pattern identification, such as the preferences of market participants. In fact, many areas have seen successful applications of AI predictions. For example, Kleinberg et al. (2017; 2018) show that AI predicted judicial decisions on bail and discrimination are more accurate and social-preferable than traditional judge-made ones. However, going from specific predicting tasks to general predictions is harder than we think. For example, IBM's Watson beat human contestants on Jeopardy but also made a number of incorrect diagnosis for cancer⁸. In addition, a key requirement for AI to work is to have a sufficiently large training dataset

⁸ See Christie (2018) "Artificial intelligence: Winter is coming" in Financial Times.

with well-structured classifications. With today's data availability and quality, however, this is not an extremely difficult task.

3.3.2 Automation

Automation is defined as the development and adoption of new technologies that enable capital to be substituted for labor in a range of tasks (Acemoglu and Restrepo 2019b). An expansive body of literature examines the effects of automation on labor supply demand, productivity and inequality (Acemoglu and Restrepo 2018, 2019a; Graetz and Michaels 2018; Autor, Levy and Murnane 2003; Goos and Manning 2007; Michaels, Natraj and Van Reenen 2014). The question we ask here is in what kind of industries can automation replace humans. The key criterion used to evaluate this replacement decision rests on whether we have enough data on a specific task on the table.

To answer this question, we need to go back to the discussions of the coordination choice between markets and firms. In markets, we compare our preferences with what is available, and then decide whether and with whom to contract. Market participants encounter decisions of this kind every time they engage in the market. Therefore, we have lots of data to train AI systems. In contrast, data generated within firms, particularly by management, is relatively limited. There are several reasons. First, executives in firms face a wide variety of different types of decisions. Decisions of a similar type are few, so the data available for a particular type of decision is limited. Second, executives face a more varied set of decision options than market participants. More data is needed for machine learning systems to work, not less. Third, data about the actual process of decision-making is still lacking. Without knowing how various inputs were evaluated, weighted, compared, and then translated into a decision, machine learning is stunted. Given the diversity of decisions to be made, it is hard to imagine machine learning systems taking over general managerial decision-making from humans anytime soon.

3.3.3 Blockchain

A blockchain is a chain of blocks that each includes a number of ledger entries (Yermack 2017). Each entry contains some information of interest. Blockchains are beneficial in three situations, where the users want to be pseudonymous; the data is unstructured (use hashes to represent the data); and the data needs to be indelible, i.e., no one can alter the data without everyone else finding out. The tools brought by blockchain, such as cryptography that protects anonymity,

hashes that represents the data⁹, and proof-of-work that verifies data, are particularly useful in these situations. With these tools, blockchains have advantages of mitigating attendant problems of falsification, double-spending, and anonymity.

An important feature of blockchains is that they can incorporate smart contracts. A smart contract is essentially a digital agreement. The main difference between smart and traditional contracts is that the clauses are not written in English and executed by lawyers. Instead, smart contracts are written in codes, with pre-programmed clauses that automatically execute themselves following a set of instructions that work on a principle of “if this happens, then execute this”. In other words, the contract is self-fulfilling and carries out what it has been coded to do. For example, it is plausible that a person’s will on the disposition of specific types of physical assets could be turned into a smart contract, with the rules of how assets should be transferred enshrined in codes. Family power plots, squabbles and lies over who gets what would be a battle with a computer program. It is important to note that the degree to which smart contracts are legally enforceable is still in question and they will not replace lawyers altogether. If the language of an insurance policy was unclear, it would still need a qualified lawyer to make a judgement.¹⁰

Blockchain has several important implications for corporate structures¹¹. For example, it is feasible to incorporate a firm without stock. Ownership and control rights to cash flow can be incorporated into smart contracts. Blockchain companies don’t need to adhere to a conventional management or financial infrastructure. They may not even need traditional bank accounts. As an example, Ethereum has raised funds through a public crowd-sale, backs a free-floating token, and operates the Ethereum Foundation as a Swiss non-profit corporation. This means that most of its value comes from sources that operate outside the legal bounds of traditional corporate structure.

Similarly, blockchain protocols (similar to traditional firms) can now be operated without contributors (similar to traditional employees) knowing each other. Blockchain protocols are the most distributed “firms” in the world. Their developers and contributors are

⁹ A blockchain contains data and information in blocks with the same size. Each block stores the hashed data from the previous block, and new data generated in the current block-forming window, to provide cryptographic security. The hashing process uses SHA256 function, which is a one way hash process. It can transform most, if not all, types of data (perfect or not) into a hash key, e.g., a series of numbers. For example, written languages, transcripts of audio, relational data, descriptions of images (arguably the most common forms of unstructured data), can be transformed using the hash function and stored on a blockchain.

¹⁰ This discussion is adapted from Buterin (2014).

¹¹ For discussions on blockchain and corporate governance, see Yermack (2017).

spread across the globe. In some cases, they are fully anonymous that the team members don't even know each other. Such anonymity can be especially desirable for employees located in more strictly regulated locations. Combined with geographical dispersion, anonymity can also have important implications for changing the written rules of corporations. For example, norms like decision-making based on seniority and other internal politics will not necessarily be intuitive, and new rules may have to be written into companies' software, rather than ingrained in their culture. Blockchain "firms" (protocols) are already thinking about solutions to help "open organisations" codify their voting, politics, and employee rewards.

Lastly, with the help of blockchains, we are approaching a state where a firm can be operated without assets. Many crypto networks still rely on physical hardware, since Bitcoin miners incur costs for electricity and computing rigs. However, newer crypto networks are already modelling themselves after the sharing economy. For example, Helium is pioneering a hardware device that provides wireless coverage. Users who buy the device can earn tokens for keeping it plugged in. Helium could effectively kickstart a decentralized telecom network, bypassing the need for capital-intensive mobile towers. Although such initiatives are still at their infant stage, they lack the burdens of real estate, labor and, sometimes, a profit motive.

4. Conclusion and future research directions

Technology is changing the way we coordinate. Developments in data ontology are helping us extract valuable data from huge streams of information and categorize it along many dimensions; progress in matching algorithms is enabling us to find and select the optimal transaction partner in the market of our choice; and machine learning systems are facilitating to identify our preferences so that we do not need to spend time to compute our multifaceted preferences and their utility-adjusted weightings. The boundaries between markets and firms are moving due to the rapid progress in technology. This is because the transaction costs incurred in both coordination methods are differentially affected by technology developments. Industries that can generate large, high quality datasets are most likely to be affected by the advances of technology.

Regarding future research directions, an area worth exploring, and closely related to our work, is how these technological developments will change the optimal size of the firm? As emphasized in Coase (1937), when discussing the nature of firm, an important factor that determines the size of the firm is management costs, which in turn, are determined by

managers' bounded rationality and firms' specific tasks. With the current technology development in automation, AI, and blockchain, management costs may decrease substantially. However, the transaction costs in the market are also reducing. As a result, whether technology increases or decreases firm size is in essence an empirical question. In addition, if technology increases (or decreases) firm size, what are the mechanisms? Does it slow the speed of diminishing return to management (change of the gradient of the marginal return and firm size curve)? Or does it reduce the coordination costs at all levels of firm size (change of level of the curve)?

The answers to these questions will have profound implications for firm regulations, product market competition, governance as well as industrial structures. Scholarly works that seek to understand the implications of each type of technology are also promising. Several many research papers examine the impact of AI, blockchains and automation on various economic and finance activities (e.g., Fisch, 2019; An, et al., 2019; Acemoglu and Restrepo 2019b; Kleinberg et al., 2017; 2018). Insights from these studies can shed important light on future applications of technology, and consequently improve societal well-being as a whole.

5. References

- Acemoglu, D., and Restrepo, P. 2018. "Low-skill and High-skill Automation." *Journal of Human Capital* 12(2): 204-232.
- Acemoglu, D., and Restrepo, P. 2019a. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*: Forthcoming.
- Acemoglu, D., and Restrepo, P. 2019b. "Automation and New Tasks: How Technology Displaces and Reinstates Labor." *Journal of Economic Perspectives* 33(2): 3-30.
- Alchian, A. A., and Demsetz, H. 1972. "Production, Information Costs, and Economic Organization." *The American economic review* 62(5): 777-795.
- An, Jiafu, Tinghua Duan, Wenxuan Hou, and Xinyu Xu. 2019. "Initial Coin Offerings and Entrepreneurial Finance: The Role of Founders' Characteristics." *Journal of Alternative Investments* 21(4): 26-40.
- Arrow, K. J. 1969. "The Organization of Economic Activity: Issues Pertinent to the Choice of Market versus Nonmarket Allocation." *The analysis and evaluation of public expenditure: the PPB system* 1: 59-73.
- Athey, Susan. 2017. "Beyond Prediction: Using Big Data for Policy Problems." *Science* 355(6324): 483-485.
- Athey, Susan. 2018. "The Impact of Machine Learning on Economics." In *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Autor, D. H., Levy, F., and Murnane, R. J. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118(4): 1279-1333.

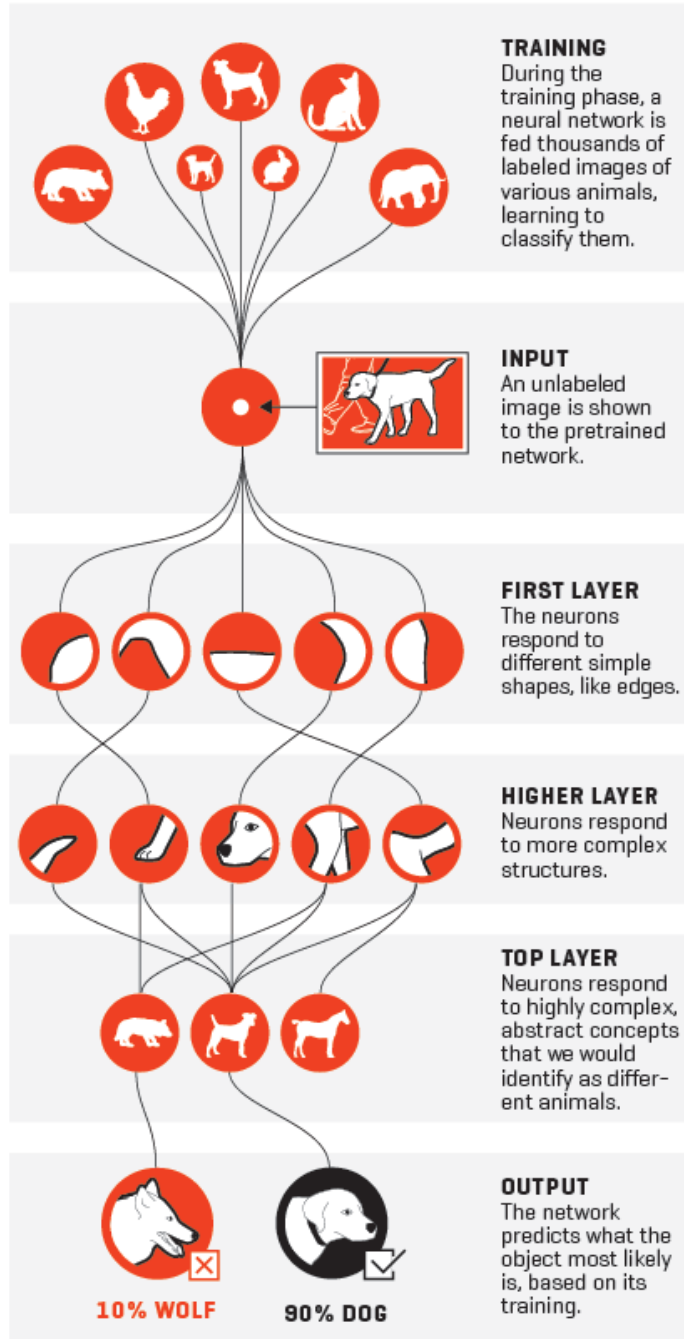
- Brynjolfsson, E., and McAfee, A. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: WW Norton and Company.
- Buterin, V. 2014. "DAOs, DACs, DAs and More: An Incomplete Terminology Guide." Ethereum Blog.
- Carlton, D. W. 1976. "Vertical Integration in Competitive Markets under Uncertainty." *Journal of Industrial Economics* 27(3): 189-209.
- Coase, R. H. 1937. "The Nature of the Firm." *Economica*, 4(16): 386-405.
- Coase, R. H. 1960. "The Problem of Social Cost." *The Journal of Law and Economics* 3: 1-44.
- Commons, J. R. 1936. "Institutional Economics." *The American Economic Review* 26(1): 237-249.
- Costa, P. T., Jr., and McCrae, R. R. 1992. *Revised NEO Personality Inventory (Neo-PI-R) and NEO Five-Factor Inventory (NEO-FFI): Professional manual*. Odessa, FL: Psychological Assessment Resources.
- Einav, L, and Levin, J. 2014. "Economics in the Age of Big Data." *Science* 346(6210): 1243089.
- Entrepreneurial Finance: The Role of Founders' Characteristics." *The Journal of Alternative*
- Farrell, J., and Saloner, G. 1988. "Coordination through Committees and Markets." *The RAND Journal of Economics* 19(2): 235-252.
- Fisch, Christian. 2019. "Initial Coin Offerings (ICOs) to Finance New Ventures." *Journal of Business Venturing* 34 (1): 1-22.
- Ford, M. 2015. *Rise of the Robots: Technology and the Threat of a Jobless Future*. New York: Basic Books.
- Goos, M., and Manning, A. 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain". *The Review of Economics and Statistics* 89(1): 118-133.
- Graetz, G., and Michaels, G. 2018. "Robots at Work." *Review of Economics and Statistics* 100(5): 753-768.
- Hayek, F. A. 1933. "The Trend of Economic Thinking." *Economica* 1(40): 121-137.
- Howell, S. T., Niessner, M., and Yermack, D. 2018. "Initial Coin Offerings: Financing Growth with Cryptocurrency Token Sales". NBER working paper, Cambridge.
- Investments* 21(4): 26-40.
- Kaldor, N. 1934. "The Equilibrium of the Firm." *The Economic Journal* 44(173): 60-76.
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., and Mullainathan, S. 2017. "Human Decisions and Machine Predictions." *The Quarterly Journal of Economics* 133(1): 237-293.
- Kleinberg, J., Ludwig, J., Mullainathan, S., and Sunstein, C. R. 2018. "Discrimination in the Age of Algorithms." *Journal of Legal Analysis*: Forthcoming.
- Knight, F. H. 2012. *Risk, uncertainty and profit*. Courier Corporation. North Chelmsford: Courier Corporation.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521(7553): 436-444.

- Michaels, G., Natraj, A., and Van Reenen, J. 2014. "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-five Years". *The Review of Economics and Statistics* 96(1): 60-77.
- Miller, G. A. 1956. "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information." *Psychological Review* 63(2): 81-97.
- Nordhaus, W. D. 2007. "Two Centuries of Productivity Growth in Computing." *The Journal of Economic History*, 67(1), 128-159.
- Noy, N. F., & McGuinness, D. L. 2001. "Ontology development 101: A guide to creating your first ontology." Working paper.
- Park, G., Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Kosinski, M., Stillwell, D. J. and Seligman, M. E. 2015. "Automatic Personality Assessment through Social Media Language." *Journal of Personality and Social Psychology* 108(6): 934-952.
- Plant, A. 1932. "Trends in Business Administration." *Economica* (35): 45-62.
- Raskin, M., and Yermack, D. 2016. "Digital Currencies, Decentralized Ledgers, and the Future of Central Banking." NBER working paper, Cambridge.
- Robbins, L. 1932. "The Nature and Significance of Economic Science." *The Philosophy of Economics: An Anthology I*: 73-99.
- Robertson, D. H. 1923. *The Control of Industry (No.4)*. Harcourt: Brace.
- Robinson, A. 1934. "The Problem of Management and the Size of Firms." *The Economic Journal* 44(174): 242-257.
- Rostek, M. J. and Yoder, N. 2018. "Matching with Complementary Contracts". Working paper.
- Usher, A. P. 1920. *An Introduction to the Industrial History of England*. Boston: Houghton Mifflin.
- Varian, H. 2018. "Artificial Intelligence, Economics and IO". NBER working paper, Cambridge.
- Varian, Hal R. 2014. "Big Data: New Tricks for Econometrics." *Journal of Economic Perspectives* 28(2): 3-28.
- Wansink, B., Kent, R. J., and Hoch, S. J. 1998. "An Anchoring and Adjustment Model of Purchase Quantity Decisions." *Journal of Marketing Research* 35(1): 71-81.
- Wansink, B., Painter, J. E., and North, J. 2005. "Bottomless Bowls: Why Visual Cues of Portion Size May Influence Intake." *Obesity research* 13(1): 93-100.
- Williamson, O. 1971. "The Vertical Integration of Production: Market Failure Considerations." *The American Economic Review* 61(2): 112-123
- Williamson, O. 1973. "Markets and Hierarchies: Some Elementary Considerations." *The American economic review* 63(2): 316-325.
- Witten, Ian H., Eibe Frank, Mark A. Hall, and Christopher J. Pal. 2016. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann.
- Yermack, D. 2017. "Corporate Governance and Blockchains." *Review of Finance* 21(1): 7-31.
- Yermack, D. 2018. "FinTech in Sub-Saharan Africa: What Has Worked Well, and What Hasn't." NBER working paper, Cambridge.

Youyou, W., Kosinski, M., and Stillwell, D. 2015. "Computer-based Personality Judgments Are More Accurate than Those Made by Humans." *Proceedings of the National Academy of Sciences* 112(4): 1036-1040.


Figure 1

HOW NEURAL NETWORKS RECOGNIZE A DOG IN A PHOTO

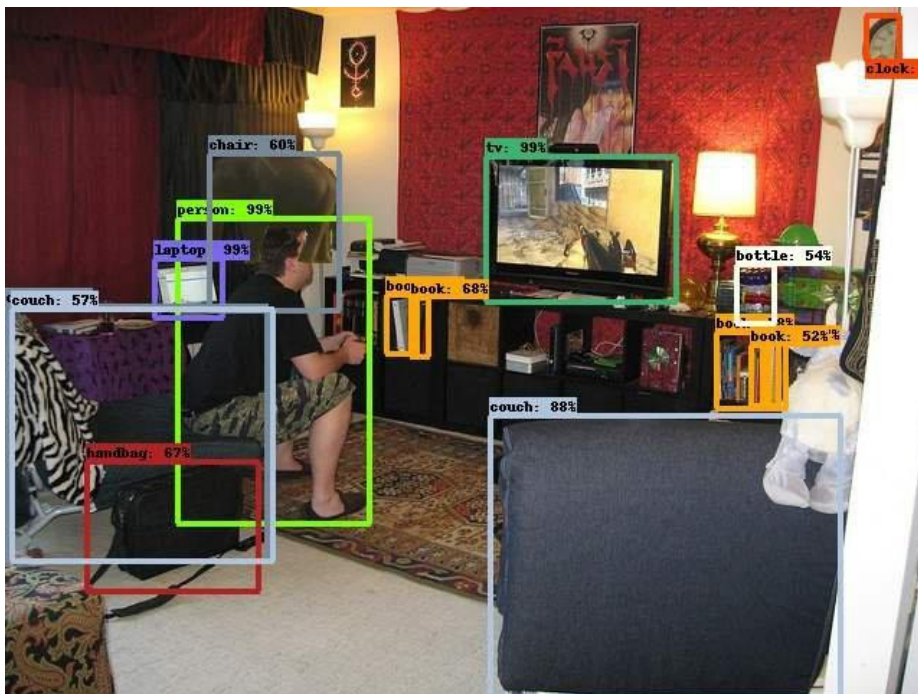


Source: <http://fortune.com/ai-artificial-intelligence-deep-machine-learning/>

Figure 2
Panel A



Panel B



Source: <https://www.quantamagazine.org/machine-learning-confronts-the-elephant-in-the-room-20180920/>