Accepted Manuscript


Title: The Brain of the Future and the Viability of Democratic Governance: The Role of Artificial Intelligence, Cognitive Machines, and Viable Systems

Author: Alexandre Perez Casaresa

PII: S0016-3287(17)30289-6
DOI: https://doi.org/10.1016/j.futures.2018.05.002
Reference: JFTR 2315

To appear in:

Received date: 22-8-2017
Revised date: 26-4-2018
Accepted date: 20-5-2018

Please cite this article as: Alexandre Perez Casaresa, The Brain of the Future and the Viability of Democratic Governance: The Role of Artificial Intelligence, Cognitive Machines, and Viable Systems, Futures https://doi.org/10.1016/j.futures.2018.05.002

# The Brain of the Future and the Viability of Democratic Governance: The Role of Artificial Intelligence, Cognitive Machines, and Viable Systems

Alexandre Perez Casares[a,*]

[a] *President, The Altius Society at Oxford, United Kingdom*

**Permanent Address:**
Alexandre Perez Casares
13 Mount Carmel Chambers
Dukes Lane
W8 4JW
London

United Kingdom



* Corresponding author.
*E-mail address:* alex_perez@hks12.harvard.edu (A. P. Casares).

ARTICLE INFO

*Keywords:*
Artificial intelligence
Viable systems
Democratic governance
Technological disruption
Income inequality
Communication systems
Complexity

ABSTRACT

At a time when smartphones, smart-homes, smart-cities, wearables, factories, etc., are becoming increasingly omnipresent, shall we also expect technological progress in artificial intelligence (AI) to result in the emergence of smart-governments and nations? The field of AI, and more broadly the development of artificial *cognitive machines*, is making breathtaking advances. A significant increase in computing power is enabling the rapid adoption of a new paradigm of AI, whereby *cognitive machines* are no longer programmed line by line, instruction by instruction, but instead are now capable of learning autonomously, thereby continuously developing themselves. The new paradigm is unleashing extraordinary progress in a wide range of applications, from healthcare to transportation and even the justice system; at the same time, these new forms of intelligence are making decisions in complex ways that escape the limits of human comprehension.

The technological transformation driven by artificial *cognitive machines* is already beginning to have far reaching consequences, some unintended, at a scale and pace which exceeds even that of the first and second industrial revolutions. This paper aims to discuss these implications, both with an ontological and epistemological perspective, to assess the potential challenges for the systemic viability of democratic societies: a fundamental change in the nature of economic wealth creation, which could raise significant social tensions, combined with the paradoxical reduction in the effectiveness of human communication. It then explores the potential role of artificial *cognitive machines* to address its own challenges, introducing the idea of a future *brain* for public governance, understanding *brain* as an emergent property resulting from the interaction of human agents and AI systems created to address the priorities of social organizations.

## 1. Introduction. The *Golden Age* of *Cognitive Machines*: a new paradigm of artificial intelligence.

In recent years, a paradigm change in the design of artificial intelligence (AI) is catalyzing an age of extraordinary progress for the development of '*Cognitive Machines*': from 'programmable computers' (Turing machine equivalent) that permitted significant algorithmic power, to 'neural networks' which enable bottom-up pattern matching capabilities. This paradigm change has been essential to address increasingly complex tasks. Resolving AI challenges like understanding language cannot be reduced to a small set of logical rules; instead, it is easier to learn statistical associations based on underlying principles of information processing. Humans themselves require years of training to absorb sufficient information to be able to converse meaningfully. The hard truth is that AI design is significantly more complex than anticipated in the mid 20$^{th}$ century: the clean symbolic world of *Good Old Fashioned AI* (Haugeland, 1985) and the reduction of human knowledge to a small set of 'expert rules' is long gone. Data driven, 'bottom-up' AI is here to stay.

Early ideas of working learning algorithms, inspired by biological neural networks that constitute animal brains, date back to the 1960s (Ivanenko, 1972). However, it was not until the last decade that these *deep learning* models become increasingly popular, as significant computing power and data were made available, a trend illustrated by Koomey's Law (Fig. 1) – extended version of Moore's Law, which describes the improvements in computation power expressed as the energy consumed to achieve certain amount of computation. According to Koomey's analysis, the number of computations per joule of energy dissipated has been doubling approximately every 1.57 years over the last seven decades (roughly equivalent to a 100-factor improvement per decade). Moreover, there could still be a significant theoretical headroom for further improvements, as the limitations imposed by Landauer's principle and the second law of thermodynamics would still leave enough margin to continue the current trend until 2048 before reaching the theoretical limits (Lambson et al, 2011).

[FIGURE IN SEPARATE FILE]

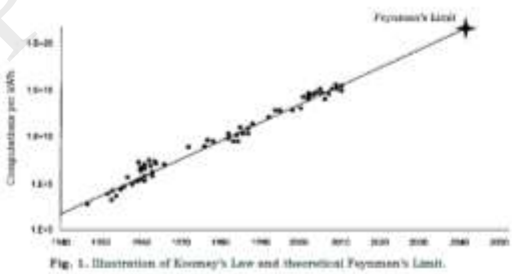

**Fig. 1.** Illustration of Koomey's Law and theoretical Feynman's Limit.

Furthermore, beyond the general increase of processing power, in the specific intersection between *deep learning* and advanced computing hardware, we are experiencing an even faster pace of development. A recent evidence is the release of the DGX-1 P100 Nvidia Deep Learning Supercomputer, with a theoretical capacity of 170 teraflops, which has achieved an improvement of over 10 times the computing power of an average *deep learning* research lab

12 months earlier, making GPU-Computing a significant performance enabler for the coming years, and a mechanism to overcome the potential limitations of Dennard scaling (Dennard et al., 1974).

Nevertheless, despite the extraordinary progress of the last decade, future steps for the development of AI architecture design shall include the ability to effectively transfer abstract learnings from one specific problem setting into a completely different situation. In 2013, DeepMind Technologies, a British company owned by Google focused in the development of AI, completed an artificial system capable of addressing a wide variety of 2D Atari games without the need of any previously pre-defined set of rules (Volodymyr et al., 2013). However, it was unable to transfer implicit knowledge between the games: every time a new type of game was played, the system would need to 'learn' how to solve it, as if it were incapable of developing any transferable intuition or generalized knowledge of the *world*.

The next frontier will be the development of transferrable, higher level, abstract intuitions. Going from perceptual AI to natural reasoning and understanding the perceived world will require the absorption of tremendous amounts of information. Even for 'perceptual' tasks such as vision, context can be vital; roughly one third of the human brain is involved in human visual processing, suggesting that similarly enormous amounts of computation will be required by any artificial system if it is to rival human visual processing and comprehension. The capacity to model the world inside the artificial agent itself is key to develop more advanced intelligence beyond pure pattern recognition: With *supervised learning*, it is possible to have a computer improving the accuracy of its neural network by performing iterations in an artificial setting. However, learning in the physical world has the inherent limitation of the *actual* speed of the world, which in computing terms is rather slow. Thus, the need to develop its own model of the world in order to achieve a successful *unsupervised learning*. Much in the same way newborn *learn* the concept of 'object', of 'gravity', etc.; they learn in a remarkably short time period, thanks to the development of its own simplified model of the *world*. Fortunately, some recent developments indicate the possibility of achieving this evolution of AI in the near future, using techniques such as neural Turing machines, NTMs, or differentiable neural computers, DNCs (Graves et al. 2014, 2016).

It is important to highlight the time horizon of the analysis presented in this paper. Predicting whether those advances will surpass human capabilities across every domain, what Bostrom denominates *artificial general intelligence* or *superintelligence* (Bostrom, 2014), is well beyond the scope of this paper. Even if the discussion presented here is forward-looking, including the analysis of future implications over the coming two or three decades, it is limited to the adoption of already existing technologies. An exercise of a different nature would be required to study a potential subsequent era, assuming such more ambitious technological predictions are materialized.

The goal of the next sections (sections 2 and 3) is to review what artificial *cognitive machines* may imply for social organizations, as a means to assess the potential impact to their viability. This assessment includes ontological and epistemological perspectives. First, a review of a fundamental issue: the potential implications for the economic organization of society, including understanding the nature of new wealth created. Then, the discussion will follow with the analysis of the impact for human perception and our understanding of knowledge, which may be severely affected by a further disruption of human communication and the flow of information.

Finally, section 4 presents a proposal to address the shortcomings diagnosed in the previous analysis. The assessment of this proposition in its capacity to improve the resilience and viability of social organizations will be understood under the framework of '*organizational cybernetics*' (Ruiz-Martin et al., 2017), which applies communication and control cybernetic principles to organizations (Beer, 1981). This discussion builds on the conceptual construction of governance and democracy introduced by Bula and Espejo (2012), which is based on the cybernetic ideas of *'complexity'* (Ashby, 1964) and *'viability'* (Beer, 1979), and underpins the comprehension of a *'democratic state'* as more than a set of formal institutions, but instead as a viable system that needs effective strategies to manage complexity under ever changing external and internal environments that constantly defy its stability. Furthermore, *Ashby's Law of Requisite Variety* tells us that a complex system requires a complex controlling device. On this regard, the power to leverage technologies of all kinds to enable adequate organizational structures to manage the increasing complexity, as discussed by Bula and Espejo, is here expanded by the potential use of *cognitive machines*. More precisely, the proposal in section 4 introduces the use of artificial *cognitive machines* dedicated to public governance, which could resolve some of the key challenges (both fundamental and informational) that are already starting to emerge in democratic societies.

## 2. The ontological perspective. Implications of the age of *Cognitive Machines* for employment, the broader economy and politics.

*Research question: "What is the nature and scale of the potential disruption for aggregated and individual economic performance, and how they may affect the viability of the system of economic organization in western democracies?"*

### 2.1. *Introduction*

It is important to analyze the processes of economic transformation in historical terms (Casares et al., 2014). There is a significant number of studies that assess the impact of major historical events for the structure of our societies. From the Neolithic Revolution, to the raise of the main monotheist religions, such as Judaism, Christianism, or Islam; the main empires, such as Macedonia or the Roman empire; significant advances in basic science, such as *calculus* or *Newtonian classical physics*; periods of significant artistic creativity, such as the Renaissance; the human colonization of major territories, such as the discovery of the Americas by Western Europeans, etc. The general conclusion of the historical analysis of these major events, is that taken in isolation, all of them were extremely significant, and had a major impact in the capacity of human civilization to be more effectively organized and generate wealth (Acemoglu et al., 2001, 2002; Engerman et al., 1997, 2002 and Diamond, 1997). However, a more recent study by Ian Morris (Morris, 2010), which analyzes the main magnitudes of social development and world human population, infers that when all those major events are contextualized under a long-term historical perspective, none of them had a very substantial impact in relative terms for global economic productivity (Fig. 2). For several millennia, per capita economic growth barely reached an annual average of 0.2 per cent, even in the most developed regions (Maddisson, 2007).

[FIGURE IN SEPARATE FILE]

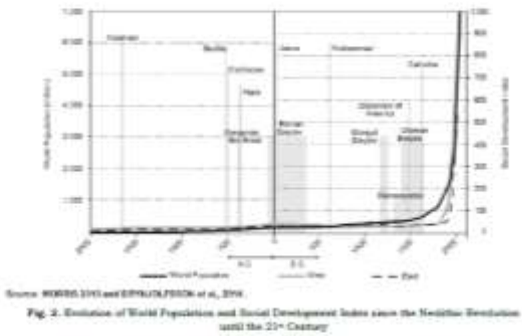

**Fig. 2.** Evolution of World Population and Social Development Index since the Neolithic Revolution until the 21st Century

However, during the 19th century a radical transformation is initiated, enabled by the development of steam power, and the ensuing industrial revolution. Borrowing the terminology from Erik Brynjolfsson, the *first machine age* (Brynjolfsson et al., 2014) was the first instance in human history when human and animal force were replaced in a significant proportion by the capacity of mechanical artificial devices; machines capable of performing physical work at a scale unimaginable until then. This innovation, together with its successive improvements and developments thanks to electric power and the internal combustion engine (ICE), catalyzed economic growth at rates close to or even above 2 per cent over decades (Gordon, 2012). This resulted in the period of fastest development experienced by human civilization.

In the second half of the 20th century, a potentially more significant transformation begins. The development of digital computation, leveraging the silicon transistor, and the advance of digital communications are this time overcoming the limitations of human cognitive capacity with the use of artificial processes, further catalyzing economic productivity growth. The analysis of annual productivity growth in the first decade of the 21st century, confirms that the United States has achieved during this period an average annual growth rate of 2.4 per cent, which even exceeds the long-term average since the start of the industrial revolution, even if recent performance may be partially offset by the impact of the *great recession* suffered since 2008 (Brynjolfsson et al., 2014).

Despite the nature of this economic transformation, some researchers question whether the modest productivity growth in the United States and the United Kingdom during the second decade of the 21st century may indicate early signs of its potential limits (Gordon, 2016). Although this will be addressed in the following section (2.2), it is worth to emphasize that the relevance of this transformation for the viability assessment of the current system of governance in western economies is not only a matter of aggregated productivity rates, but more importantly, the different nature of newly created wealth and the potential inequality of its distribution (sections 2.3 and 2.4).

Nevertheless, this process illustrates the development of a new model of economic productivity, where human-designed devices enable further productivity beyond the limitations of human cognitive capabilities. This new context may be called the '*Age of Cognitive Machines'*.

As introduced in the previous section, the main technological innovations that are catalyzing this transformation include four main domains: a new paradigm of artificial intelligence architecture, the fast development of digital computation capacity, increased digital connectivity and the use of advanced robotics. These four technological advances may be summarized in the ability of artificial machines to process information faster and more efficiently, leveraging significant volumes of data, and enabled by the growing usage of digital communication infrastructure.

The implications of this process expand well beyond the digital economy, reaching a wide range of products, services and industrial sectors. If the transformation initiated during the 19th century and early 20th century was unique in human history, then the economy of *cognitive machines* has further accelerated this process since the second half of the 20st century. The rise of *cognitive machines* may remain the most relevant economic driver for a significant portion of the 21st century.

Beyond the opportunities posed by the direct impact of increased productivity, the economy of *cognitive machines* has major effects in four fundamental aspects of social organization: i) the nature of new wealth creation and the complexity of its measurement, ii) the role of human labor in the economy, iii) the dynamics of income and wealth inequality, and iv) the scale and asymmetry of international interdependence. These four factors significantly increase the level of economic complexity.

### 2.2. *The Nature of New Wealth Creation*

New wealth creation is quite different from previous historical periods in three dimensions. First, the radical decrease of marginal production costs as a result of technological transformation. The pace of cost reduction has led some economists to even question the path and sustainability of recent economic growth. Jeremy Rifkin is one of the strongest advocates of this idea and coined the concept of a *zero marginal cost society* (Rifkin, 2014). Rifkin argues that the digital nature of new goods and services results in deflationary processes with severe consequences that threaten the stability of the market-based economic system. He states that the increase in productivity with very low marginal costs may drive prices down significantly, threatening the mechanisms of investment and wealth creation, which are the foundation of capitalism. However, a more holistic analysis indicates that the aforementioned deflationist process is often compensated by an even larger increase in consumption volume, resulting in net positive economic growth. This process is well illustrated in the example of the computational transistor. While the cost of a transistor has decreased over 50 per cent per year, it has been more than offset by an extraordinary increase in demand. The combined effect of both dynamics is a net annual increase of 18 per cent in economic contribution of computing technologies over the last five decades. This is a qualitative transformation, beyond a sequence of incremental improvements, which in turn enables new applications completely unthinkable in the original scenario of unitary costs, driving historically high growth rates in economic productivity in the long run.

The second characteristic is the accelerated path of transformation. The extended Moore's law is just one example of the phenomena of sustained accelerated growth, with 100 years of future growth equivalent in magnitude to the previous 20,000. Unlike the first and second industrial revolutions, or the British agricultural revolution which took place over multiple generations, the age of *cognitive machines* and its subsequent disruption of multiple industries is substantially faster.

The third characteristic is the increasingly complex accounting of total economic welfare, which may lead to underestimate the impact in gross domestic product and human development. This is of particular importance, as the widespread focus on conventional statistics of economic productivity has led several researchers to conclude that we are in a period of significant economic underperformance (Gordon, 2016). Economists and statisticians acknowledge that the framework for economic measurement needs to evolve as a result of the increasing weight of digital goods and knowledge production (Haskel et al., 2017; Corrado et al., 2017 and Diewert et al., 2017). Indeed, a common trait of the four technological trends mentioned above, is that the process of digitization of production has a relatively low degree of dependence from physical products and processes, *exclusive and non-rival goods* in economic terms (DeLong, 2015). Thus, increasing the complexity to quantify the economic impact, in particular its effect in gross domestic product accounting beyond marketed economic activity.

An illustration of this accounting paradox is the evolution of the industry of music production and distribution. Fig. 3. illustrates the evolution of per capita consumption of recorded music, which reached its historic peak at the beginning of the last decade in the United States. Since then, observed data *a priori* suggests a dramatic reduction of consumption without precedents. However, the actual evolution of recorded music consumption has been the opposite: at the end of the decade of the 2000s, the number of hours of recorded music consumed per capita reached its historic peak, while the economic contribution to *gross domestic product* –and, hence, to accounted *productivity*– was at its minimum. In other words, we 'consume' more music than ever, although in digital format, however this consumption seems to have only a minor impact in the traditional magnitude of economic production: GDP. This phenomenon, and the limitation of techniques used in traditional economic theory, has urged many economists to underestimate the potential for economic growth and prosperity of this new era (Cowen, 2011; Gordon, 2016).

In macroeconomic terms, the new economy of *Cognitive Machines* generates the appearance of a moderate contribution to economic welfare, when indeed we are experiencing a significant transformation in productivity and 'consumption'. If DeLong's analysis is correct, we are currently undergoing a period of historic welfare growth, beyond observed figures in traditional metrics of national accounts. Private corporations are producing more goods than ever before, and the population is in turn consuming record per capita levels of goods and services, both physical and digital, from music and entertainment to garments and transport, at significantly lower cost. Nevertheless, despite the seemingly idyllic situation, it also generates major challenges which will be addressed in the following sections (2.3 and 2.4).

[FIGURE IN SEPARATE FILE]

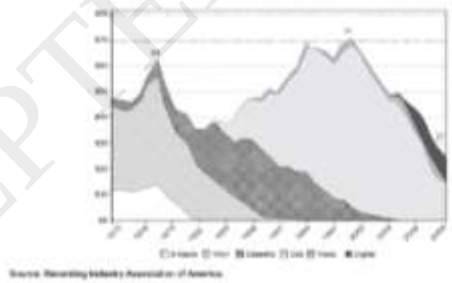

**Fig. 3.** Evolution of per-capita consumption of recorded music in the United States (2010 US$)

### 2.3. *The role of 'human labor' in the age of Cognitive Machines: 'Technological unemployment'?*

The discussion about the impact of fast technological disruption for labor markets and employment is not new. Groups of English textile workers and weavers in the 19th century, known as *luddites*, rebelled against the impact that new industrial machines would have over their employment, as these machines had the potential to make the time spent learning the skills of their craft worthless. In 1930, John M. Keynes presented a more optimistic vision of technology in his essay about the '*economic possibilities for our grandchildren'* (Keynes, 1930). Even if Keynes highlights the potential positive impact of innovation and technology for economic prosperity, he also acknowledges a concern about *technological employment*, as new methods of production would be discovered with a pace and scale larger than the ability to find new types of occupation for human workers.

The mainstream view of modern economics is that most innovation and industrialization processes, at least those experienced until now, have not eliminated the need for human labor. Indeed, the opposite: this progress has enabled enough new opportunities to satisfy the growing population during the 19th and 20th centuries. However, the question is whether the most recent economic transformation since the late 20th century, and more so in the 21st century based in the advancement of *Cognitive Machines*, could depart from previous observations. For the first time in over two centuries, the strong increase in productivity in developed economies is not directly '*trickling down*' into growing wages, in real terms (Fig. 4.), a phenomenon that has been coupled with one of the lowest employment participation rates of the last decades (Fig. 5.), even in countries such as the United States where the rate of unemployment is historically low.

[FIGURE IN SEPARATE FILE]

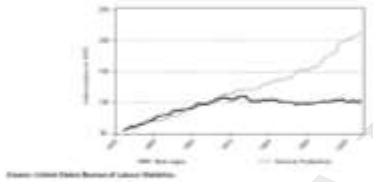

**Fig. 4.** Evolution of Economic Productivity Growth and Real Wages in the United States, 1945-2010

[FIGURE IN SEPARATE FILE]

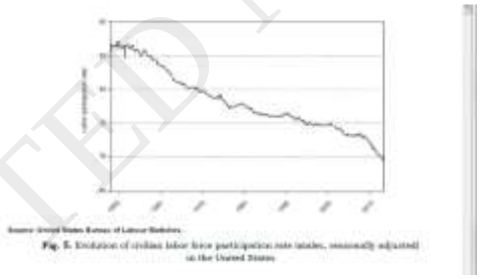

**Fig. 5.** Evolution of civilian labor force participation rate (males, seasonally adjusted) in the United States

Furthermore, this trend could potentially continue and accelerate further over the coming decades: simulations performed by Carl B. Frey and Michael Osborne at the University of Oxford, based on current available technologies applied to a diverse range of human labor tasks, indicate that over 47% of current occupations in the United States have a high risk of being significantly disrupted by machines over the coming two decades (Frey et al., 2013). A more recent study by the Organization for Economic Co-operation and Development (OECD), based on a more disaggregated occupational classification (Nedelkoska et al., 2018), estimates that 14 per cent of jobs in developed countries are "highly automatable", and a further 32 per cent will likely experience significant change in the way they are performed. Both studies emphasize the vast magnitude of the potential disruption, despite being based exclusively on what is feasible with current technological capabilities. Even if no further breakthroughs occur, the mere adoption of current automation technologies would result in the fastest rate of labor disruption ever recorded. In other words, no major human civilization has ever experienced such a significant disruption in employment, both in magnitude and timeframe. This is not to say that we will definitely experience a significant degree of unemployment, but at least an extraordinary reallocation of human occupations, and a major transformation of how these and other occupations are performed. This trend will stress the current model of economic growth and the viability of the social market economy of most major economies, and in turn the stability of liberal democracies.

### 2.4. *Viability of the free market economic model, the social contract, and the current system of governance in western democracies.*

Each of the innovation processes mentioned above share a common characteristic: its differentiated impact over each of the three main socio-economic stakeholders. For the purpose of this analysis, we may simplify the diversity of human stakeholders in society to 3 main type of agents according to their role in the economic system (note that these groups are not mutually exclusive):

A. The "*innovators*": direct creators and contributors to a disruptive technology. These receive not only labor income, but significant capital gains (i.e., total or partial ownership of the share capital of the new technology). Capital gains frequently outweigh the amount of employment cash income, as a result of the increasingly global nature of technological adoption (Mankiw, 2013). The combination of a high degree of globalization and the heavy share of 'non-physical assets' of many new technologies catalyzes this trend: highly educated employees and entrepreneurs are able to capture extraordinary incomes at a scale that was unthinkable decades ago. Erik Brynjolfsson and Andrew McAfee (Brynjolfsson et al., 2011) highlight: "*Aided by digital technologies, entrepreneurs, CEOs, entertainment stars, and financial executives have been able to leverage their talents across global markets and capture reward that would have been unimaginable in earlier times.*" An example of this dynamic is the case of Travis Kalanick and the value of his personal shareholding ownership in Uber, the company that he founded in 2009, currently estimated at over US$5 billion.

B. Second, the "*consumers*" who enjoy the benefits of a new technological development. As an illustrative example, users of a piece of software called Turbotax may enjoy a service that facilitates the preparation of their tax returns, a complex endeavor in the United States, at a significantly lower price (near $50) compared with traditional human tax advisors, and frequently with a higher degree of precision in most general circumstances.

C. And finally, the "*displaced workers*", whose occupations are directly or indirectly disrupted, and in some cases completely replaced by the new *cognitive machines.* Building on the Turbotax example, Brynjolfsson (2014) highlights that human tax consultants in the United States have lost a significant portion of their income over the last few years as a result of the competition from the new Turbotax software.

The growing presence and power of *Cognitive Machines* has a diverging impact on each of these three groups of stakeholders. The most immediate consequence is the rapid increase of inequality beyond historical records, with a particular negative effect for middle and low-income families in developed economies, whose jobs are most susceptible to be disrupted by new technological developments (Frey et al. 2013). Fig. 6. indicates the enormous growth of the share of income concentrated in top households in the U.S. ("top-5" and "top-20" percentiles) over the last few decades, while middle and low-income households suffered a significant reduction of their share of income over the same time period.

[FIGURE IN SEPARATE FILE]

**Fig. 6.** Change in Share of Total Income in the United States by percentile, 1967-2012

There are diverse explanations for the growing inequality in developed economies. The approach that is presented above departs partially from the thesis of Tomas Piketty (2014). Piketty's grand theory of capital and inequality, where other things being equal, slower economic growth increases the importance of wealth in a society, drives him to conclude that only a burst of rapid growth, from technological progress or rising population, or government intervention can prevent economies from returning to "patrimonial capitalism." More precisely, Piketty argues that the origin of a significant portion of the observed inequality is related to the different rates of return of capital and labor, and therefore his critique is often focused on a system of hereditary elites that monopolize the control over capital. However, the perspective presented here establishes instead that the origin of the increasing inequality is linked –in recent years– to the nature of new wealth creation derived from technological innovation, disconnected from any '*caste*' system, and exposed to significant volatility and further disruption over time. It is indeed the combination of a system of significant returns for those that innovate –versus those that simply inherit–, with '*winner takes all*' market dynamics (Mankiw, 2013). This has significant implications for the future of capitalism and the design of public policy, specifically the need to reassess the role of the state and a new social contract (Soros, 1997); this includes the debate around the potential necessity of some form of universal basic income, the role of public ownership of innovation capital (Rodrik, 2015), the future of the education system in order to augment (rather than replace) human abilities by leveraging *cognitive machines*, etc.

In summary, the previous discussion highlights the magnitude and speed of change in the economic system as a consequence of the new economy of *cognitive machines*, and the significant stress that it will pose to the current systems of governance in western democracies. This is driven by the combination of the increasing challenge to properly measure welfare and other economic statistics, the scale of disruption to the labor market, and the unequal distribution of benefits of this economic transformation.

While it is true that other large-scale economic transformations have taken place in the past, the current rate of disruption is now faster than any of those. By comparison, an equivalent disruption to 46 per cent of the workforce[2] –estimates from the comprehensive study by Nedelkoska et al. (2018)– took place over the course of more than half a century in the United Kingdom during the 19th century. In other words, major historical technology adoption cycles in the past were even longer than the average worker productive lifetime. This is a fundamental difference, as the current disruption is now taking place in less than a third of the average working life, while at the same time key institutional elements of democratic governance, such as the education system, have not yet been properly reinforced or adapted with this challenge in mind.

Regardless of the final precise magnitude of its impact, what is clear is that the increasing rate of change driven by the widespread adoption of *cognitive machines* will pose a growing stress to the current systems of governance in western democracies, and it demands organizational design changes to be implemented in order to preserve their systemic viability. The increasing complexity in the governance of change is further enhanced by the fact that these systems of governance were mostly designed at a time when the rate of change was an order of magnitude slower, taking place over generations, and thus the current potential risk

[2] Note that in this context, the expression "*disruption to the workforce*" is frequently misunderstood as employment losses. Indeed, it may not necessarily imply job displacement at all. The estimates include the share of occupations that will be disrupted to a degree in which most of the skills required today to perform those activities will no longer be needed, thus requiring a significant reskilling of those workers performing the affected occupations.

for the viability of the current system, in cybernetic terms. As a result, and referring to Ashby's *Law of Requisite Variety*, the increasing complexity in the environment affecting the organization requires a corresponding increase in *variety* within the control system, which underpins the need to introduce certain improvements to the western democratic system of governance, as discussed in section 4. Solving the actual economic challenges will require changes in methods and institutions far beyond Economic Policy. The changes must be institutional and cannot be delayed, stressing that these are not so much improvements in Economic Policy as they are changes in Political Economy.

## 3. The epistemological problem of perception in the age of *cognitive machines*. When technology increases the gap between truth and belief.

*Research question: Beyond the fundamental issues that our economies will face in the coming years as a result of the extended use of artificial intelligence systems, which were discussed in the previous section, is there a risk that addressing these issues may not be prioritized in the policy agenda as a result of a failure of the information systems in democracy? How will 'cognitive machines' precisely affect these information systems?*

For much of the 20th century, stylized facts about the impact of technology for communication in democratic societies presumed a positive correlation between advancements in digital technologies and a more informed society and therefore robust democratic process. Indeed, as recently as 2013, it was argued that the Arab Spring represented the peak of this democratizing trend enabled by new media (Khondker, 2011 and Howard et al., 2013).

A few years later, we see evidences of the opposite. After a first wave of increasing connectivity with the adoption of digital communication technologies, the unintended side effects of artificial intelligence for the dissemination of truth begin to be widespread. We should consider that among the 50 top-performing *news-feed* stories of 2016 on Facebook (Silverman, 2016) was not only the fabricated story that "*President Obama had banned reciting the pledge of allegiance in US schools,*" with over 2,177,000 engagement actions (shares, comments and reactions) and over 20 million unique people reached, but also the story headlined: "*Pope Francis Shocks World, Endorses Donald Trump for President, Releases Statement.*" We are living at a time in which, as Matthew d'Ancona explains, we seem to prioritize the *"visceral over the rational, the deceptively simple over the honestly complex"* (d'Ancona, 2017). However, the existence of human irrationality is definitely not a novel idea, as Daniel Kahneman highlighted in his theory of the two human modes of thinking (Kahneman, 2011). Therefore, why are these phenomena recently becoming more apparent?

*The Genie in the Bottle: Artificial Intelligence and Communication in Society.*

In 2016, Oxford Dictionaries chose *post-truth* as its word of the year, defining it as shorthand for "*circumstances in which objective facts are less influential in shaping public opinion than appeals to emotion and personal belief*". But why is society becoming more reactive to non-factual pieces of information? The answer can be found in the design of the system: the increasing use of artificial intelligence to determine the inner workings of the main social digital media platforms. Most users of Facebook, Twitter, YouTube, etc., do not often realize that the *news-feed* they see is basically an artificial intelligence that is just working tirelessly to keep them engaged as much as possible, so by interacting with Facebook they are indeed interacting with an artificial intelligence system.

This has far reaching consequences. The generalized conception towards technology in modern societies is to attribute a neutral morality (Verbeek, 2011). Using the simple example of a knife, the common view would be that it is merely a tool; the knife is not good or evil *per se*, it is just humans that can use it for good or evil purposes. However, in the case of *artificial cognitive machines*, we may find unintended consequences and emergent complex dynamics that vastly exceed our human capacity to anticipate *ex ante*.

The issue with an autonomous artificial intelligence system is that when given a goal, it will precisely try to achieve it. Facebook has designed its complicated AI system with the goal of keeping the users engaged, and thus maximize its advertising revenues. The very reason why AI is used is because the tasks that the AI system is performing are too complicated for humans to deal with on their own. It is indeed the system that Facebook, as a private agent, has designed in order to deal with the increased variety (complexity). It would be hard to imagine Facebook hiring millions of editors to optimize each user experience, picking the right stories that the user may be most likely to engage with. Obviously, this would be completely unfeasible, so Facebook turns to an artificial system to achieve its goal. The problem is that while the original goal might be morally neutral, or even positive, the ways of achieving this goal are decided by the artificial system on its own, with potential consequences exceeding the human capacity to foresee. It is hard for any human to really understand how AI has developed the specific means to achieve its original goal, as its complexity quickly exceeds the limits of human comprehension. The issue is that those means may be ethically questionable, or even completely undesirable. Suddenly the AI system at Facebook may learn that showing questionable stories may be the best way to achieve the goal that was initially set by its designers, the goal of keeping the users engaged, even if that entails leveraging sophisticated psychological appeals to Kahneman's system one (based on our intuitive, non-rational processing of information). Months later, we are somehow surprised about the major epidemic of questionable *fake news* inundating most social digital platforms. This disbelief includes the designers of the system itself, as Mark Zuckerberg, Facebook's founder and chief executive, testified in front of the United States Congress in April 2018.

The fundamental elements behind the algorithm, whether it is based on neural networks or other machine learning techniques, are usually rather simple (frequently composed of simple operations such as multiplications or additions). However, it is impossible for a human mind to comprehend the emerging complexity: how starting from those very simple rules in the intelligence model, the system may evolve when millions of elements are progressively processed. The end result of this ecology is a situation in which individuals and societies are making decisions based on what an artificial system has told them, without having a real understanding of why this decision has been made. The developers of the AI system at Facebook understand the simple rules that they request the machine to follow, rules that ultimately determine which stories and news are shown to the user in his *newsfeed*. However, the exact intricacies of this machine learning model that is actually driving choices, quickly escapes their comprehension.

*Empirical evidence during the 2016 presidential election in the United States: a matter of communication disruption by AI, not of economic transformation (yet).*

The magnitude of the economic transformation enabled by the age of *cognitive machines* discussed in the previous section will likely characterize the political discourse in future electoral processes. However, there is increasing empirical evidence to conclude that these economic effects were still not the focus of the recent political events, even in the case of the United States. Instead, data suggests that among other factors it was education, and not income, that predicted who would vote for Donald Trump: according to Pew Research Center, the main gap in share of votes by demographic groups was a 39-points difference among 'white voters with low-basic education' (i.e., 39 percentage points higher share of votes for Trump vs.

Clinton). It is important to highlight that despite the correlation between income and education, the average Trump supporter had higher income than the average Clinton voter ($72,000 vs. $62,000 average household income). Furthermore, only 33% of Trump supporters considered the 'gap between the rich and poor' a 'very important problem', compared to 72% among Clinton supporters. Indeed, only two other issues scored an even larger gap between Trump and Clinton voters: 'gun violence' (31% vs. 73%), and 'climate change' (15% vs. 66%), with lower relevance among Trump supporters in all cases.

It is also paradoxical that the economic policies introduced by the Trump administration not only fail to address the potential economic impact of artificial *cognitive machines* on its voters, but the opposite. One example of such policies is the proposed healthcare reform, the American Health Care Act, AHCA, known colloquially as '*repeal and replace'*, initially approved by the United States House of Representatives. According to the Congresional Budget Office (2017), it is estimated that AHCA would result in a loss of coverage for 23 million people over the course of the next decade, while cutting taxes by about $765 billion, of which the top 1 percent of income earners would receive 40.2% of the tax savings (Urban-Brookings, 2017).

Nevertheless, the first explicit effect of *cognitive machines* was its extraordinary impact for the effectiveness of political communication. Despite the inscrutability of most social digital communication platforms in regards to their data, there is statistical evidence available to illustrate the dynamics discussed above. Not only the most popular news stories widely shared in Facebook during the 2016 election were factually incorrect and tended to favor Donald Trump over Hillary Clinton (Silverman, 2016), but also the level of activity of political *artificial bots* reached a historic record (Narayanan et al, 2018 and Kollanyl et al., 2016), with a significant bias towards pro-Trump content (4.9 pro-Trump tweets from highly automated accounts per each pro-Clinton tweet in the 8 days before election day). Furthermore, data suggests that most users seem to share news online without ever reading them, but still impacting the readers of the headline (Gabielkov et al. 2016).

Critics of this view argue that the actual risks posed by technology for the emergence of *echo chambers* are overstated, given the increasing proliferation of countless online media outlets catalyzed by the growing use of digital media platforms. Indeed, empirical evidence demonstrates that, *a priori*, a high-choice media environment leads individuals to more diverse content and perspectives, preventing them from finding themselves in an *echo chamber* (Dubois, 2018). However, these studies miss the most significant trend affecting media content consumption in new digital devices over the very last few years: the increasing reliance on the algorithmic system of content filtering behind the most commonly used platforms, such as Facebook or Twitter, which in practice is eliminating that choice. Nevertheless, the aim of this paper is not to precisely quantify the magnitude of this interference, which is still a pending task despite recent research efforts, but instead to highlight the potential risk which demands the pertinent organizational design changes to preserve the ability of information systems to deal with the increased *variety* (in cybernetic terms) and in turn, ensure the resilience and viability of the democratic system.

## 4. Conclusion. A proposal for a resilient democratic society in the age of *cognitive machines:* the necessary coexistence of private AI agents and public AI systems of governance.

"*The moral to be drawn from this dangerous nightmare situation is a simple one: don't let it happen*," said George Orwell when he tried to clarify the meaning of his dystopian novel *1984*, in which he imagined a society in thrall to a diabolical overseer, wherein even the truths of mathematics could be replaced by "*alternative facts*".

In broad terms, the ever-increasing volume of data combined with the complexity of the new challenges that derive from the use of artificial *cognitive machines* itself, are exceeding the current capacity of most social agents, as well as their information systems. This situation is incompatible with the requirements for the viability of democratic societies, as detailed in the *viable system model,* VSM[3] (Beer, 1979). According to this model, in order for the system to survive (i.e., to be viable), it must be adaptable, among other requirements, and thus organized in such a way as to meet the demands of surviving in the changing environment, as otherwise the pathologies of the organization will threaten is viability. This has implications both for the specific elements of the system, as well as its communication channels (including the insufficient presence of algedonic channels). In the example of the criminal judicial system which will be detailed in the following paragraphs, the judges, as components of the *operational system*, are unable to fulfill their goals effectively. At the same time, the issues in the communication channels highlighted in the previous section hinder the effectiveness of the current democratic process, by restricting the ability of citizens to perform informed decisions when electing its democratic representatives, indicating a pathology of the information systems and communication channels. Furthermore, the lack of supranational coordination required to address the emergent international asymmetry of these issues, is a sign of a structural pathology according to organizational cybernetics, which states the need for different levels of recursion to address the complexity that emerges at each level (Perez Rios, 2012). A more detailed analysis of the specific considerations for each element of the system shall be performed applying this and other models of organizational cybernetics (Espejo et al., 2011, and Schwaninger, 2008).

The paradoxical conclusion of the combination of both the ontological and epistemological analysis of the age of *cognitive machines*, is that the use of artificial systems of intelligence for social governance (Fig. 7) may be precisely a necessary mechanism to address the increased complexity and the emerging organizational pathologies which result from the extended use of artificial machines by private agents –of course, this assumes a desire to preserve a democratic system of social organization, as alternative models of governance are beyond the scope of this paper.– The combination of these multiple artificial systems of intelligence (in collaboration with human agents) will progressively result in the future emergence of a collective *brain* for public governance: the *brain* of the social system.

The power to leverage technologies of all kinds to enable adequate organizational structures as discussed by Bula and Espejo (2012), is therefore here expanded by the potential use of *cognitive machines*, which facilitate the necessary increase in *variety* within the control system.

[3] The VSM establishes the necessary and sufficient conditions for the viability of an organization. These are related to the existence in an organization of a set of systems or functions (named S1 through S5), as well as a set of relationships among these and the environment. Communication channels are responsible for connecting all those systems or functions, as well as linking the organization with its environment. Particularly relevant are the *algedonic channels*, whose role is collecting and transmitting information critical for the viability of the organization.
Failures in the organizations' design or in how it operates may be associated to a kind of pathology. Pérez-Ríos (2010) classified the most frequent ones in three main groups: 1) structural pathologies, related to organizations' structural design, and to how the organization copes with its total environmental complexity by creating the necessary sub-organizations, 2) functional pathologies, related to the adequacy of the organizations (at all recursion levels) to the prescription made by the VSM about functional subsystems and their relations, and 3) information system and communication channels pathologies.

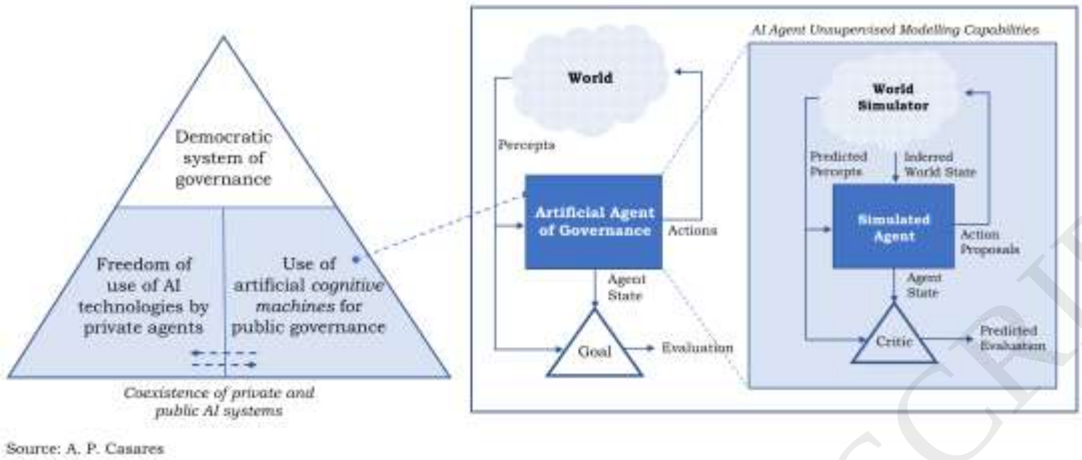


**Fig. 7.** Artificial cognitive machines with unsupervised learning capabilities are necessary agents to sustain the viability of a democratic system of governance, assuming the freedom of use of AI systems by private agents is preserved.

Several economists and computer scientists have already made significant advances in Experiential Economics, with soft agents, growing models of viable institutions (Agent Based Modeling) in the new field of *Artificial Economics* (Hernández, 2017) and Viable and Resilient Systems (Ruiz-Martin et al., 2017). These developments underpin the potential to include in our experimental designs of institutions, new agents with cognitive endowments (*cognitive machines*). Progressively, *Artificial Economics* may become standard methods to support modeling physical systems populated by social agents (humans and or artificial cognitive machines) and may bring significant advances in Economics and Social Sciences in what we may call the '*Illustration of the 21st century.*'

Early examples of the practical implementation of coexisting human and artificial agents in public governance already start to emerge in domains with an intrinsic conservative tradition, such as the criminal judicial system. As of 2017, courts and corrections departments around the United States already use algorithms to determine a defendant's "risk", which ranges from the probability that an individual will commit another crime to the likelihood a defendant will appear for his or her court date. These algorithmic outputs inform decisions about bail, sentencing, and parole. Judges are required by law to base their decision solely on the prediction about the expected behavior of the defendant, making this an ideal application of artificial *cognitive machines* for public decision making.

Each AI tool aspires to improve on the accuracy of human decision-making that allows for a better allocation of finite resources. At first, it may be hard to be comforted by the idea of a judge that while she has the ultimate decision power to slam his wooden hammer, she is in fact relying on a computer screen that says, for example, that it has evaluated this particular defendant should not be granted parole because of his chances of reoffending being too high, a prediction result of its artificial learning over the analysis of hundreds of variables and millions of previous cases. However, recent research indicates a significant precision superiority of these artificial systems vs. human judges (Kleinberg et al., 2017), estimating a potential of over 40% reduction in jail population awaiting trial if the target crime rate by defendants is kept unchanged. In the absence of such an AI system, a judge that may have seen the defendant for only a few minutes, would then have just a few seconds to skim through his case data and millions of previous case materials to make his decision.

The overwhelming increase of available information and case complexity makes this task unfeasible for a human brain. However, the process followed by the AI system is deterministic, and it is applied to all previous cases. Moreover, the evidence seems to demonstrate that it works effectively for the benefit of humans: reducing crime rates while limiting jail times. No doubt, relevant questions arise about the moral implications of following decisions without a clear understanding of how they are made. Furthermore, judges are now increasingly avoiding to disagree with the recommendations made by the artificial program, namely because no judge would like to take responsibility for shortening the sentence against the advice of the AI system, and then having the defendant reoffend. As a result, in practice, we now have the cases of freedom and imprisonment decided by an artificial system, precisely because humans do not completely understand how the intelligence of these systems works as to contradict their judgments. Also, empirical evidence appears to demonstrate that humans are not necessarily better judges and are also subject to other limitations: an Afro-American defendant with the same criminal history compared to a white defendant will get a longer sentence by a human-only judge vs. when assisted by the AI system referred above. As the volume of available information and the complexity of cases continues to increase, the use of AI systems –in collaboration with human agents– for the public organization of society is emerging as an effective and economically efficient mechanism to increase the viability of our societal goals.

Nonetheless, it may be useful to explore human communications to check the effective application of justice. Indeed, if the organizational structure underpinning the justice system is inadequate, no doubt judges will be better off applying the AI algorithm, but if the structure was more effective, it might be possible to complement effectively the AI system. This underpins the importance of leveraging AI to *augment* (rather than just *replace*) human intelligence, and the effectiveness of public administration.

While a key contribution of this paper is to highlight the necessary coexistence of private AI agents with artificial systems for public governance decision making, multiple open questions about this new ecology require additional exploration. As anticipated, there are two areas of particular concern that would require further inquiry: (i) a discussion of the ethical and moral considerations of the use of such a system of governance; (ii) the persistence of the proposed analysis in the long-run, particularly if the continuous advance of artificial *cognitive machines* expands beyond a potential theoretical inflection point (Bostrom, 2014 and Barrat, 2015), in which a machine's *brain* would be capable of general intelligence action.

**References**

ACEMOGLU, D.; JOHNSON, S. and ROBINSON, J. (2001). «The Colonial Origins of Comparative Development: An Empirical Investigation», *American Economic Review* 91, pp. 1369-1401.

ACEMOGLU, D.; JOHNSON, S. and ROBINSON, J. (2002). «Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution». *Quarterly Journal of Economics* 117, pp. 1231-1294.

ASHBY, R. (1964). *An Introduction to Cybernetics*, Methuen & Co, London.

BARRAT, J. (2015). *Our Final Invention: Artificial Intelligence and the End of the Human Era*, 1st edition, Nueva York. Thomas Dunne Books, St. Martin's Griffin.

BEER, S. (1979). *The Heart of Enterprise*, Wiley, Chichester.

BEER, S. (1981). *Brain of the Firm*, 2nd edition, Wiley, Chichester.

BOSTROM, N. (2014). *Superintelligence: Paths, Dangers, Strategies*, 1st ed., Oxford University Press, Oxford.

BULA, G. and ESPEJO, R. (2012). «Governance and inclusive democracy», *Kybernetes*, Vol. 41, No. 3/4 pp. 339-347.

BRYNJOLFSSON, E. and MCAFEE, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, 1st edition, W. W. Norton & Company, Nueva York.

CASARES, A. P. and MUÑIZ, M. (2014). «Progreso tecnológico y orden internacional: hacia una nueva economía y una mejor gobernanza», *Información Comercial Española, Revista de Economía*, No. 880, pp. 39-55, October.

Congressional Budget Office (2017), *Cost Estimate for H.R. 1628, American Health Care Act of 2017*, www.cbo.gov/publication/52752.

CORRADO, C., HASKEL, J. and JONA-LASINIO, C. (2017). «Knowledge spillovers, ICT and productivity growth». *Oxford Bulletin of Economics and Statistics*, *79*(4), pp.592-618.

COWEN, T. (2011). *The Great Stagnation: How America Ate All the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better*, 1st edition, Penguin Group, Nueva York.

D'ANCONA, M. (2017). *Post-Truth. The New War on Truth and How to Fight Back.* Random House.

DELONG, B. (2015). «Making Do With More», *Project Syndicate, The World's Opinion Page*, February, Prague.

DENNARD, R. H., GAENSSLEN, F., YU, H., RIDEOUT, L., BASSOUS, E. and LEBLANC, A. (1974). «Design of ion-implanted MOSFET's with very small physical dimensions», *IEEE Journal of Solid State Circuits*. Volume: 9, Issue: 5.

DIAMOND, J. (1997). *Guns, Germs, and Steel: The Fates of Human Societies,* New York, W.W. Norton & Company.

DIEWERT, E., and FOX, K. (2017). «The Digital Economy, GDP and Consumer Welfare.» Presented at the NBER Summer Institute CRIW Workshop 17-18 July 2017.

DUBOIS, E., and BLANK, G. (2018). «The echo chamber is overstated: the moderating effect of political interest and diverse media», *Information, Communication & Society*, Volume: 21, 2018, pp.729-745.

ENGERMAN, S. and SOKOLOFF, K. (1997). «Factor Endowments, Institutions, and Differential Paths of Growth Among New World Economies: a View from Economic Historians of the United States» in: HARBER, STEPHEN (Ed.), *How Latin America Fell Behind*. Stanford University Press, Stanford, pp. 260-304.

ENGERMAN, S. and SOKOLOFF, K. (2002). «Factor Endowments, Inequality, and Paths of Development Among New World Economies», *National Bureau of Economic Research*, Working Paper 9259.

ESPEJO, R. and REYES, A. (2011). *Organizational Systems: Managing Complexity with the Viable System Model*, Springer, New York.

FOREIGN AFFAIRS (2015). BEZOS, J.; BESSEN J.; CHRISTENSEN, et al. «Here Come the Disrupters!» *Foreign Affairs*. Volume 94, Number 1.

FREY, C. B. and OSBORNE, M. (2013). «The Future of Employment: How Susceptible are Jobs to Computerization», Oxford Martin School Publications, September, Oxford, United Kingdom.

GABIELKOV, M., RAMACHANDRAN, A., CHAINTREAU, A., & LEGOUT, A. (2016). Social clicks: What and who gets read on twitter? In *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science* (pp. 179-192). ACM.

GRAVES, A.; WAYNE, G. and DANIHELKA, I. (2014). *Neural Turing machines*. Preprint at http://arxiv.org/abs/1410.5401

GRAVES, A and WAYNE, G. (2016). Hybrid computing using a neural network with dynamic external memory. *Nature* 538, 471–476

GORDON, R. J. (2012). «Is US economy growth over? Faltering innovation confronts the six headwinds», *National Bureau of Economic Research*. Working Paper nº 18315, August, Cambridge, Massachusetts.

GORDON, R. J. (2016). *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*, Princeton, Princeton University Press, ISBN 978-0691147727.

HAUGELAND, J. (1985). *Artificial Intelligence: The Very Idea*, Cambridge, Mass: MIT Press, ISBN 0-262-08153-9.

HASKEL, J.; CORRADO, C.; FOX, K.; GOODRIDGE, P.; JONA-LASINIO, C.; SICHEL, D. and WESTLAKE, S. (2017). «Improving GDP: Demolishing, Repointing or Extending?» Indigo Prize Entry.

HERNANDEZ, C. (2017). «Sixty Years of Economics: Some Lessons for the Future». In: HERNANDEZ, C. (eds) *Advances in Management Engineering. Lecture Notes in Management and Industrial Engineering.* Springer, Cham. 95-107

HOWARD, P. N., and HUSSAIN, M. M. (2013). *Democracy's fourth wave? digital media and the Arab Spring,* Oxford University Press on Demand.

IVAKHNENKO, A. G., and LAPA, V. G. (1967). *Cybernetics and forecasting techniques*. American Elsevier Pub. Co.

KAHNEMAN, D. (2011). *Thinking, Fast and Slow*, 1st edition, Macmillan, New York.

KEYNES, J. M. (1930). «Economic Possibilities for our Grandchildren», *Essays in Persuasion*, W.W. Norton & Co., Nueva York, pp. 358-373.
KHONDKER, H. H. (2011). Role of the new media in the Arab Spring. *Globalizations*, 8(5), 675-679.
KLEINBERG, J., LAKKARAJU, H., LESKOVEC, J., LUDWIG, J., and MULLAINATHAN, S. (2017). Human decisions and machine predictions (No. w23180). *National Bureau of Economic Research.* Chicago
KOLLANYL, B., HOWARD, P., and WOOLLEY, S. (2016). «Bots and Automation over Twitter during the U.S. Election.» Data Memo 2016.4, Project on Computational Propaganda, Oxford.
KURZWEIL, R. (2012). *How to Create a Mind: The Secret of Human Thought Revealed*, Duckworth Overlook, London.
LAMBSON, B., CARLTON, D. and BOKOR, J. (2011). «Exploring the Thermodynamic Limits of Computation in Integrated Systems: Magnetic Memory, Nanomagnetic Logic, and the Landauer Limit», *Physical Review Letters*, vol. 107, no. 1, Document No. 010604.
MADDISON, A. (2007). *Contours of the World Economy 1-2030 AD: Essays in Macro-Economic History*, 2007 edition, Oxford University Press, Oxford.
MANKIW, N. G. (2013). «Defending the One Percent», *Journal of Economic Perspectives*—Volume 27, Number 3, pp. 21–34.
MATTEI, U., and RUSSI, L. (2012). The evil technology hypothesis: a deep ecological reading of international law. *Cardozo L. Rev. DeNovo 263.*
MNIH, V., KAVUKCUOGLU, K., SILVER, D., GRAVES, A., ANTONOGLOU, I., WIERSTRA, D. and RIEDMILLER, M. (2013). «Playing Atari with Deep Reinforcement Learning». arXiv:1312.5602
MORRIS, I. (2010). *Why the West rules – for now. The cycles of history and what they tell us about the future*, 1st edition, Farrar, Straus and Giroux, Nueva York.
NARAYANAN, V., BARASH, V., KELLY, J., KOLLANYI, B., NEUDERT, L. and HOWARD, P. N. (2018). «Polarization, Partisanship and Junk News Consumption over Social Media in the US». Data Memo 2018.1. *Project on Computational Propaganda.* Oxford.
NEDELKOSKA, L. and QUINTINI, G. (2018). «Automation, skills use and training». *OECD Social, Employment and Migration Working Papers*, No. 202, OECD Publishing, Paris.
PEREZ RIOS, J. (2010). «Models of organizational cybernetics for diagnosis and design», *Kybernetes*, Vol. 39 Issue: 9/10, pp.1529-1550.
PEREZ RIOS, J. (2012). Design and Diagnosis for Sustainable Organizations. The Viable System Method. Springer, Heidelberg, New York, London.
PIKETTY. T. (2014) *Capital in the Twenty-First Century*, 1st edition, Harvard University Press, Cambridge.
RIFKIN, J. (2014). *The Zero Marginal Cost Society. The Internet of Things, The Collaborative Commons, and The Eclipse of Capitalism*, 1st edition, Palgrave Macmillan, Nueva York.
RODRIK, D. (2015). «From Welfare State to Innovation State». Project Syndicate, January, Princeton.
RUIZ-MARTIN, C., PEREZ RIOS, J.M., WAINER, G., PAJARES, J., HERNANDEZ, C., LOPEZ-PAREDES, A. (2017). «The Application of the Viable System Model to Enhance Organizational Resilience», in: HERNANDEZ, C. (eds) *Advances in Management Engineering. Lecture Notes in Management and Industrial Engineering.* Springer, Cham, pp. 95-107.
SCHWANINGER, M. (2006). *Intelligent organizations: Powerful models for systemic management*, Springer Berlin, Heidelberg.
SILVERMAN, C. (2016). «This Analysis Shows How Fake Election News Stories Outperformed Real News on Facebook». *BuzzFeed News.*
SOROS, G. (1997). «The Capitalist Threat». *The Atlantic*, January.
THIEL, P. (2014). *Zero to One: Notes on Startups or How to Build the Future,* Virgin Books, London.
Urban-Brookings Tax Policy Center (2017). *Distribution of Major Tax Provisions in the American Health Care Act of 2017*, version 0217-1. Washington.
VERBEEK, P. (2011). *Moralizing Technology: Understanding and Designing the Morality of Things.* University of Chicago Press, Chicago.

* * *