COLLEGIUM OF ECONOMIC ANALYSIS
WORKING PAPER SERIES

The Digital Era, Viewed From a Perspective of Millennia of Economic Growth

Jakub Growiec

SGH KAE Working Papers Series Number: 2018/034 April 2018

The Digital Era, Viewed From a Perspective of Millennia of Economic Growth*

Jakub Growiec†

March 29, 2018

Abstract

I propose a synthetic theory of economic growth and technological progress over the entire human history. Based on this theory as well as on the analogies with three previous eras (the hunter-gatherer era, the agricultural era and the industrial era) and the technological revolutions which initiated them, I draw conclusions for the contemporary digital era. I argue that each opening of a new era adds a new, previously inactive dimension of economic development, and redefines the key inputs and output of the production process. Economic growth accelerates across the consecutive eras, but there are also big shifts in factor shares and inequality. The two key inputs to the digital-era production process are hardware and software. Human skilled labor is complementary to hardware and substitutable with software, which increasingly includes sophisticated artificial intelligence (AI) technologies. I also argue that economists have not yet designed sufficient measurement tools, economic policies and institutions appropriate for the digital-era economy.

Keywords: economic growth, technological progress, unified growth theory, digital economy, artificial intelligence.

JEL codes: O10, O30, O40.

This is a paper about the future of global economic growth and technological progress. Therefore it extensively deals with the past.

*Helpful suggestions by Katarzyna Growiec and Ingmar Schumacher are gratefully acknowledged. The views expressed in this paper belong to the author only, and have not been endorsed by NBP, SGH, or any other institution. I solemnly apologize for any possible mistakes. If you see a factual error, please let me know. If you think that my theoretical postulates are *kind of right* but not *quite right*, and you would like to suggest modifications and extensions, that's all I wish for.

†SGH Warsaw School of Economics and Narodowy Bank Polski, jakub.growiec@sgh.waw.pl.

1 Introduction

The primary objective of this paper is to propose a synthetic theory of economic growth and technological progress over the entire human history, with an intention to frame the millennia of past development as a single cumulative process. Inspired by the excellent books by Diamond (1997), an anthropologist, ecologist and geographer, and Harari (2014), a historian, I organize the millennia of human history into four eras of economic development: the hunter-gatherer era, the agricultural era, the industrial era and the digital era. The eras have been initiated by four respective revolutions: (i) the cognitive (Upper Paleolithic) revolution (ca. 70 000 BP) – the time when *homo sapiens* acquired its status of a dominant hominin species and began gaining control over natural habitats across the world; (ii) the agricultural (Neolithic) revolution (ca. 10 000 BP), marking the beginning of a transition from hunting and gathering to sedentary agriculture; (iii) the industrial revolution (ca. 1800 CE), the onset of industrial production, capital accumulation, and systematic increases in GDP per capita; and (iv) the digital revolution (ca. 1980 CE), marking a rapid explosion in the world's capacity to compute, store, and communicate data. I emphasize the co-existence of the eras in time.

Across all the eras I consider the following themes: knowledge accumulation, economic growth, key factors of production and their mutual relation, inequality, and side effects of development.

As a theoretical brace connecting all four eras, I propose that the driving forces of development can be summarized in a single encompassing concept of *local control*. The drive to maximize local control is an emergent feature of human behavior, which appears regardless of our final goals, provided only that we value the preservation of ourselves and our children.

The second objective of this paper is to draw new conclusions for the digital era, building on the synthetic theory of economic growth and technological progress, presented in this paper, as well as the analogies with previous eras and the technological revolutions which separated them. I am able to obtain new useful results here thanks to taking a broader perspective on the digital age than it has been typically done in the literature. I argue that looking further back in time and investigating the consecutive technological and economic revolutions provides new insights for the digital era, compared to viewing it as a mere continuation of the industrial era.

The motivation for writing this paper is the growing dissonance between the empirically documented properties of the digital economy and the way in which it is usually analyzed in the macroeconomic literature, both empirically: e.g., growth and development accounting based on National Accounts data (Jorgenson and Stiroh, 2000; Timmer and van Ark, 2005; Jorgenson, 2005; Fernald, 2015; Gordon, 2016), and theoretically, in formal models of economic growth and technological change, e.g. relating to the form of the R&D equation (Jones, 1999; Ha and Howitt, 2007;

Madsen, 2008; Bloom, Jones, Van Reenen, and Webb, 2017; Kruse-Andersen, 2017). I think that important lessons can be learned from comparing the recent digital revolution to the industrial revolution, whose workings have been thoroughly studied in the unified growth theory literature (Hansen and Prescott, 2002; Galor, 2005, 2011), and the more ancient Neolithic (agricultural) revolution (Diamond, 1997; Hibbs and Olsson, 2004; Bocquet-Appel, 2011; Harari, 2014).

The key take-away from my reading of this heterogeneous literature is that each technological revolution is accompanied by a fundamental paradigm shift (Olsson, 2000, 2005), opening an entirely new, previously absent *dimension of economic development*, uncovering new trade-offs, and ultimately rendering the previous-era measure of development obsolete. For example, the industrial revolution initiated systematic increases in GDP per capita and growth in demand for skilled labor, opened the children quantity–quality tradeoff, and by lowering fertility broke the Malthusian link between development and population size (Galor, 2011), thus making population size an obsolete measure of economic development. In my view, it is not sufficiently appreciated by the economists that the digital revolution in fact caused equally fundamental changes in the economy. Therefore some of the current discussions, on the risks of entering a period of "secular stagnation" (Gordon, 2016), declining labor share (Karabarbounis and Neiman, 2014), increasing profit share (Barkai, 2017), increasing income inequality (Piketty, 2014), and so on, may potentially be affected by the fact that we often view the rapidly expanding digital economy through the lens of inadequate, industrial-era measurements and theories.

There are a few important conclusions from this paper. First, each opening of a new era marks the arrival of a new dimension of economic development, which records at least an order of magnitude higher growth rates than the previous-era one. Hence, this new dimension soon becomes dominant, eventually leading to a "secular stagnation" in the previous-era measure of economic development, even though the development process has in fact accelerated. However, there are sizable feedback loop effects across eras as the new era tends to initially reinforce the previous-era economy. Applying this logic to the transition between the industrial and the digital era, one should expect first a massive boost in GDP per capita (the preferred measure of economic development in the industrial era), driven by the reinforcements from the digital era, but eventually probably a "secular stagnation" when the potential for global GDP growth will be saturated, and further growth will be concentrated exclusively in the digital sphere.

Second, economists have not yet developed sufficiently detailed and reliable methods of measuring aggregate inputs (broadly categorized as "hardware" and "software") and output (tentatively summarized here as *flows of useful data*) in the production process of the digital economy. We tend to look at the digital economy through the lens of methodologies such as the National Accounts, which are well suited to the analysis of the industrial economy but obsolete when applied to the digital economy. Furthermore, being fixated on country-level analysis and policy, economists also tend to underestimate that the digital economy is fully global. Beyond economics, however, information theorists have already produced first estimates of the world's capacity to store, communicate and compute information (Hilbert and López, 2011; Gillings, Hilbert, and Kemp, 2016). These will be probably useful as a starting point for setting up more detailed digital-era accounts.

Third, our landmark economic policies and institutions have been designed with the industrial-era economy in mind. Unfortunately, they are typically either inefficient when applied to a digital-era economy, or not applicable at all. In effect the digital economy is largely unregulated. This may be contributing to a range of unfavorable outcomes such as the spread of misinformation over the Internet, formation of global monopolies, or systematic intrusions into privacy. More broadly, they are also likely contributing to the increasing (top) income inequality and a declining labor share (Karabarbounis and Neiman, 2014; Jones and Kim, 2017). The situation is akin to the early years of the industrial revolution, when only agricultural-era regulations and policies were in place, which – analogously – allowed rapid concentration of wealth, very high income inequality and low labor shares (Piketty and Zucman, 2014), until adequate industrial-era institutions were installed (and, in retrospect, until the moment when returns to human capital began to match or exceed those of physical capital).

Fourth, the two key inputs to the digital-era production process can be summarized as "hardware" and "software", with human skilled labor being complementary to "hardware" and substitutable with "software". This means that if we want to keep our jobs, we must particularly carefully trace the development of the latter. Indeed, automation is already gradually eliminating routine jobs, both manual and cognitive (Acemoglu and Autor, 2011; Frey and Osborne, 2013), and the remaining jobs are safe only until the development of sufficiently sophisticated and versatile AI technologies. Finally, the ultimate piece of "software" which may arrive in the course of the digital era, artificial general intelligence (AGI), is a double-edged sword whose consequences may vary from an utopian depiction of a technological "singularity" (Kurzweil, 2005) to human extinction (Muehlhauser and Salamon, 2012). I expect that more and more sophisticated, multi-purpose AIs, ultimately leading to an AGI, will almost surely be developed because doing so is very much in line with the human drive to maximize local control, and economic incentives are particularly huge given that software, as compared to hardware, is a clear development bottleneck of the contemporary digital-era economy.

The main limitation of the current paper is that it is qualitative. It provides an encompassing theoretical framework for studying economic growth and technological progress over the millennia of human history and into the digital era, but it does not offer any formal model which could be rigorously taken to statistical data. Its merits consist mostly in connecting the known facts, organizing them in order to build a consistent narrative and a blueprint of a unified growth theory for the entire humankind, and proposing bold but risky conjectures for the digital era, based on extrapolations and intuitive analogies to the past.

The structure of the paper is as follows. Section 2 organizes the millennia of human history into four eras, initiated by four revolutions. In Section 3 I found my theory on the proposition that maximization of human control has been a key driving force of development, active across all four eras. Section 4 is devoted to the topic of accumulation of knowledge. Section 5 formalizes the idea that each new era opens up a new dimension of economic development, documents the subsequent accelerations in economic growth following the technological revolutions, and addresses the measurement problem. Section 6 views the production processes in the consecutive eras through the lens of an aggregate production function and pinpoints the most important production factors of each era. It also makes the case for better digital policy. Section 7 discusses the side effects of development, paying special attention to the existential risk from AGI. Section 8 concludes.

2 Four Eras of Development, Initiated By Four Revolutions

The narrative of long-run economic growth and technological progress over the millennia of human history can be intuitively organized as four eras, initiated by four consecutive revolutions, such that each new era puts development into a higher gear. Let me begin with (Y-chromosomal) Adam and (mitochondrial) Eve.

Hominins first appeared on the face of Earth about 2 400 000 years ago, in East Africa (Dunsworth, 2010). According to the (arguably most popular) Out-of-Africa theory, the same geographic region was also the cradle of the anatomically modern human, *homo sapiens* (Ashraf and Galor, 2013). First humans are estimated to have appeared around 200 000 years before present (BP) and made several early efforts to expand their habitats towards Asia and Europe.

The first turning point in human history was the *cognitive revolution*, also called the Upper Paleolithic revolution. Its dating is uncertain; for orientation let us place it around 70 000 BP (Harari, 2014). Newly acquired cognitive skills, such as the theory of mind and the ability to create, document and communicate stories, gossip, tales, legends, and abstract ideas, allowed humans to become "behaviorally modern" and, crucially, advance from the middle to the top of the food chain (Tattersall, 2009; Tomasello, 2014; Harari, 2014). What followed was the unique wave of human dispersal out of Africa which turned out permanent from today's perspective – although the expansion was probably originated by as few as 1000 individuals

(Liu, Prugnolle, Manica, and Balloux, 2006). It followed a "southern route" via the Bab al-Mandab strait towards Asia and Australia, which was reached about 65 000 BP–50 000 BP. Europe, in contrast, was colonized by the *homo sapiens* only about 45 000 BP, despite geographical proximity and lack of physical barriers. The hypothesized reason for this delay is that at the time, Europe was populated by the *homo neanderthalensis*. This coincidence signifies that Europe's Neanderthals were "an additional ecological barrier for modern humans, who could only enter Europe when the demise of Neanderthals had already started" (Liu, Martinón-Torres, Cai, Xing, Tong, Pei, Sier, Wu, Edwards, Cheng, Li, Yang, Bermúdez de Castro, and Wu, 2015).

The *hunter-gatherer era* constituted the first major step of development achieved by the humankind which has never been achieved by any other species. The perspective of other species elucidates how impressive it was that hunter-gatherer bands and tribes managed to colonize almost all habitats in the world, including passing the Bering Strait to the Americas and inhabiting most Pacific archipelagos. While they were taking successive natural habitats under their control, they drove all other hominins to extinction, and did the same for a wide variety of other species such as the Asian mammoths, Australian megafauna including huge diprotodons (Harari, 2014), or American saber-toothed cats.

However, from today's perspective, they only managed to obtain miniscule "economic growth" rates. Around 5 000 BCE, the combined world population was in the order of just 5 million (Kremer, 1993), and it took the *homo sapiens* a whopping 20–25 millennia to defeat the Neanderthals and colonize Europe. Yet, this pace of development was already at least an order of magnitude faster compared to the pace of species evolution.

The second turning point in human history was the *Neolithic revolution*, also called the agricultural revolution. Domestication of plants and animals allowed for dramatic increases in the amount of calories derived from a given land area, leading to marked increases in human population density, and encouraging the formerly roving hunter-gatherer tribes to abandon their former lifestyle in favor of sedentary agriculture (Diamond, 1997; Hibbs and Olsson, 2004).

The *agricultural era* was initiated ca. 10 000 BP (8 000 BCE) in the Fertile Crescent (in the Middle East). The first eight "Neolithic founder crops" included wheat, barley, lentil, pea, and chickpea. Independent agricultural revolutions, with different founder crops, appeared subsequently over later millennia in various other parts of the world: Central China, the New Guinea Highlands, Central Mexico, northwestern South America, Sub-Saharan Africa, and eastern North America. Adoption of agriculture was only gradual because gains in food production were initially modest and technological progress as well as diffusion was slow. For these reasons, agricultural societies co-existed with roving hunter-gatherer tribes for entire millennia before the former managed to outnumber and thus economically dominate the latter (Diamond, 1997). In fact, some populations never managed to shift to sedentary agriculture, e.g. the Aboriginal Australians before the European conquest.

The third turning point in human history was the *industrial revolution*. It began around 1800 CE, first in England, and then it quickly spread across Europe and into the Western Offshoots such as the US. Formerly agricultural societies began to build factories, manufactures, and engage in mass production. The set of key technological inventions which fueled this shift included, among many others, steam engine, electricity, internal combustion engines, and indoor plumbing (Gordon, 2016). The industrial revolution was enabled by the earlier scientific revolution (including empiricism, experimentation, mathematization, and hypothesis testing), embraced by the Western societies around the 16th century.

As opposed to the agricultural revolution, the industrial revolution happened only once in history. At the time, the world was already connected, and the gap between technological knowledge acquired by the Western societies and any other disconnected society was too large to enable an independent industrial revolution.

The technology and economy of the *industrial era* were gradually adopted across the world, increasing both economic power and standards of living, and enabling acceleration of economic growth rates by an order of magnitude: doubling times were cut from hundreds of years to just decades (Galor, 2005). The transition to the industrial era is not yet complete, though: subsistence agriculture still employs a majority of population in, e.g., Sub-Saharan Africa or India. Major pockets of agricultural poverty can also be found across, e.g., China, Indonesia, the Middle East, or Latin America.

Before the fruit of the industrial era have been fully reaped, the humankind entered yet another technological revolution: the *digital revolution*. A tentative timing would place this revolution around the 1980s, when personal computers have begun to permeate firms and households; however the revolution really gained momentum in the 2000s when the Internet connected the computers in a truly global World Wide Web.

The key inventions which fueled the digital revolution were the Turing machine, semi-conductors, integrated circuits, followed by what now is the hallmark of the *digital era*: personal computers, the Internet, cell phones, and industrial robots. First key developments were achieved in the US, but in this case technology diffusion is really fast, so that nowadays a major part of the hardware frontier of the digital economy has moved to East Asia (Japan, Korea, Taiwan, China), and on the consumer side, even very poor and infrastructurally disadvantaged parts of the world frequently use (at least some) digital technologies. An illustrative statistic describing the pace of development in the digital era is Moore's Law (Hilbert and López, 2011; Bloom, Jones, Van Reenen, and Webb, 2017): since the 1980s, world's general-purpose computing capacity doubles every 1.5 years.

Looking from 2018 into the future, we cannot exclude the possibility of further technological revolutions. As argued by some authors, such a revolution may likely occur already within the current century (Hanson, 2000; Kurzweil, 2005; Brynjolfsson and McAfee, 2014), with totally unpredictable consequences.

Three observations stand out from this narrative, linking the prehistoric hunter-gatherer times with the modern digital world. First, each new era accelerates "economic growth" by at least an order of magnitude. Second, each new era opens up an entirely new avenue of development, thus redefining what "economic growth" refers to. Third, arrival of a new era does not end the previous era, and instead the technologies from the consecutive eras coincide in time. In fact, as I will argue later, each new era dramatically strongly feeds back on the previous-era economy: just think of the gains from mechanization and specialization of agriculture, or the computerization and robotization of industry.

Focusing only on the population which resides in a given place on the world map, one may get the impression that the humankind developed linearly, gradually passing from the hunter-gatherer to agricultural, to industrial, and to the digital era. However, two elements which distort this view are that (i) the timing of take-off to the next era has varied largely across the world (Diamond, 1997; Galor, 2005), and that (ii) the take-off was only gradual. Therefore the technologies of the consecutive eras co-exist not only globally (giving rise to cross-country inequality) but also locally (contributing to within-country inequality). I argue that between-era inequality is typically far greater than within-era inequality.

3 The Driving Force of Development: Maximization of Local Control

Looking back at the four technological revolutions, each of which was an impressive leap of human capability, begs the question: why did mankind achieve all this? What was the driving force of all these developments?

In the beginning, we were just lucky. Having tested plenty of other genetic designs of species, the evolutionary process produced the *homo sapiens* who slightly exceeded other hominin species in terms of frontal cortex capacity of the brain. This slight advantage turned out sufficient for the human to go and rule the world.

What I mean by being "lucky" is that evolution is not intentional when it explores the space of species designs and maximizes genetic fitness of the resultant species (commonly known as "survival of the fittest"). The evolutionary search is multidimensional: species tend to fit to certain ecological niches by developing unique advantages. It is also extremely slow: the fitness of innovations (mutations) is evaluated only *ex post* and over hundreds of generations. For large mammals like us, this means that major useful innovations are embodied in us at the time scale of hundreds of thousands of years. Hence, although species evolution never stops, we are anatomically almost (if not completely) indistinguishable from our ancestors who inhabited the Ethiopian savanna 100 000 years ago.

In consequence, evolution can only explain the *cognitive revolution* of the Upper Paleolithic but cannot explain any of the later developments of humankind. So why did they appear?

In the following subsections I put forward a theory which views the process of human expansion, economic growth, and technological progress, as a single process where humans intentionally maximize their *local control*. I view human local control maximization as one of many sub-routines of the grand process of species evolution, but the only one which got out of hand because of its sheer pace (faster by orders of magnitude than the pace of evolution), recursive self-improvement, and greed for resources. To my knowledge, the idea to use local control maximization as a brace connecting all four eras is new to the literature.

3.1 The Concept of Local Control

I postulate that what brought us to the digital era of 2018 and beyond, is a process of intentional local control maximization. This process operates on top of species evolution, dominating it in terms of speed – in the early hunter-gatherer era only slightly, but now massively.

By *maximizing local control* I mean exploring the space of possible actions in order to reflect one's preferences as closely as possible, and satisfy as many needs as possible, using the resources available in one's environment. Thus defined "local control" is not new, but I use this specific label to emphasize the versatility of the concept, its ability to describe actions and incentives across all millennia of human history, as well as the limitations of human perception and cognition. The word "control" signifies that human decisions often go beyond the economics textbook problem of maximizing utility from consumption (and possibly leisure), and that the maximized objectives may include a much wider variety of variables and motives. The word "local", in turn, is meant to emphasize that our information sets are usually limited, both in space, so that we take only part of our surroundings (natural environment, other people) into account, and in time, as uncertainty with respect to future developments forces us to be myopic.

If we accept the proposition that economic growth and technological progress have been driven over the millennia by the actions of thousands, millions, and now billions of individual humans who try to maximize their local control, it becomes natural to measure aggregate world development with aggregate measures of the extent of human local control. A caveat is that human actions are usually uncoordinated and our interests often collide. Therefore an appropriate empirical measure of aggregate human local control must be calculated on a net basis, excluding the extent of human control which is at the expense of other humans. A more thorough discussion of measurement issues is provided in Section 5.

Modern physics and computer science literature offers additional characterizations of the ultimate driving force of economic and technological development, interpreting it as *cumulative optimization power* (Yudkowsky's stance in the Hanson and Yudkowsky (2013) debate) and *choice entropy* (Wissner-Gross and Freer, 2013). For reasons that will be clear in Section 5, however, it is much easier for me to relate to *local control* than any of these two alternatives.

3.2 Why Do Humans Maximize Local Control?

Why do we maximize local control? Evolution equipped us, just like any other species, with the intention to satisfy our needs, avoid risks, and safely pass our genes to the next generation. The key difference is that the evolutionary success of all non-human species is constrained by other species and particular properties of the ecosystem; in contrast, humans have proven to be extraordinarily successful in their survival and multiplication strategies. As the only species in Earth's history, we managed to dominate the entire planet, modifying most of the world's ecosystems and driving other species to mass extinction (Kolbert, 2014).

The humankind did all this because we were never fine with making simple choices from a predefined set of alternatives. Instead, we actively seeked to expand the set of open choices (the effectiveness drive), as well as insure and diversify against risks (the safety drive, cf. Bowlby, 1969). Hence, we have been gradually loosening the constraints binding our decisions and getting around them by inventing new dimensions of action. It has been argued that the key anatomic feature, which enabled this unprecedented success, is the frontal cortex of our brains. Once its capacity has surpassed a certain threshold, the *homo sapiens* acquired the theory of mind, and used it to create, document and share stories, gossip, tales, abstract ideas, and ultimately knowledge and technology (Tomasello, 2014; Harari, 2014). This led to a unique level of versatility and adaptivity, inaccessible to any other species.

It also allowed us to organize ourselves into bands, tribes, local communities, and ultimately societies and nations. The human brain can naturally accommodate social contacts up to ca. 150 acquaintances (Dunbar, 1992, 1993). This "Dunbar's number" determined the size of early hunter-gatherer tribes, and nowadays sets the maximum size of an organization before it requires a hierarchical structure. For all other species, this number is much lower.

The beauty of the local control maximization process lies in its emergent character. The individual humans do not have to consciously follow this objective; some of us may even actively oppose it, declaring instead the pursuit of happiness, harmonious family life, etc., or recalling religious or spiritual motivations. The process emerges even if there is no ultimate human goal (no objective *meaning of life*) whatsoever or if our declared goals are mutually contradictory. The necessary condition is only that we want ourselves and our children to survive.

The emergent character of the human local control maximization process is easiest to explain by an analogy to machine intelligence, whose actions arise in the course of maximizing a predefined objective function. The instrumental convergence thesis (Omohundro, 2008; Bostrom, 2012) signifies that *regardless of the final goal* (of an AI), there is going to be convergence of its auxiliary goals, which are in fact dimensions of local control. Omohundro (2008) names the following goals: self-preservation, efficiency, resource acquisition, and creativity. Analogously, Bostrom (2012) mentions: self-preservation, goal-content integrity, cognitive enhancement, technological perfection, and resource acquisition. All that is generally dear to us humans, isn't it?

The human local control maximization process outstripped species evolution in terms of pace following the cognitive revolution. Nowadays, in the industrial and digital eras, outcomes of virtually all our innovative decisions are evaluated within a single person's lifetime, and often on a scale of months or days (if not hours or minutes). In further contrast to the evolutionary process, *ex ante* evaluation of innovations is possible and often practiced, especially when dealing with policy decisions. However, evaluation is often limited in space and time, so that we hardly ever think about the whole mankind or the infinite time horizon. The negative consequences of this fact will be reviewed in Section 7.

Finally, let me also compare the *local control* concept with the concept of *intelligence*, which can be defined, amonst other definitions, as *efficient cross-domain optimization*, an ability to hit narrow targets in broad search spaces, cf. Yudkowsky (2013). A very similar definition could be coined for local control, only without the word "efficient". This is the key difference: intelligence is about obtaining maximum possible control with minimal possible resources, whereas local control is an extensive measure which does not require efficiency in resource usage. Many goals can be achieved either by throwing in plenty of resources (the brute force solution), or by acting more intelligently, with fewer resources.

3.3 Local Control and the Hierarchy of Needs

While our innate instincts may be rooted in the reality of a hunter-gatherer tribe in prehistoric Ethiopia, our modern lives are very different. This is because the extent of control over the environment which the humankind has acquired over the millennia allows us to satisfy many more human needs today. The gradual advancement observed across the four eras can be illustrated with Maslow's hierarchy of needs (Maslow, 1954).

Hunter-gatherers can only satisfy their most fundamental needs. They spend a large chunk of their time searching for food, a *physiological* need. Furthermore, as their rudimentary equipment, which they carry with themselves, provides them with some warmth and shelter and they live in extended families as well as organize into bands and tribes, we may say that some of their basic *safety* and *social belonging* needs are also satisfied.

Compared to hunter-gatherers, sedentary subsistence farmers can satisfy somewhat more *safety* needs thanks to their permanent dwellings, food storage, elimination of many threats (e.g., predators, hostile tribes) from the ecosystem, and the accumulation of more durable goods. Following the increases in population density and growth of cities and states, *social belonging* needs can be potentially better satisfied, too.

The industrial era opened up a wide range of new possibilities for satisfying human needs. Sustained growth in GDP per capita allowed a wide group of people to benefit from increased standards of living and financial security. Developments in medicine dramatically improved human health and increased our longevity. All this can be classified as satisfaction of additional *safety* needs. The industrial economy also greatly widened the group of people who could develop their specific skills and gain recognition for their work, thus satisfying *esteem* and *self-actualization* needs.

Compared to the industrial era, the digital era offers many more avenues for satisfying *esteem* and *self-actualization* needs. The digital economy offers high rewards for creativity, innovativeness, and unique skills. The Internet further democratizes access to platforms of self-presentation, and facilitates recognition of one's work. Finally, whether we like it or not, gradual automation of routine jobs pushes us away from the most boring occupations towards more skilled and creative work.

Our systematic advancement up Maslow's (1954) hierarchy of needs is an indication of success in following Omohundro's (2008) and Bostrom's (2012) emergent drives. Satisfying low-level needs (physiological, safety, social belonging) follows from the goal of self-preservation. The drive for resource acquisition helps satisfy both low-level needs (in particular, safety) and the high-level need of esteem. Note that the resources we accumulate may vary from an inventory of consumption goods and financial wealth to social contacts, political power, and useful data. The drives toward creativity (cognitive enhancement) and efficiency (technological perfection) are, in turn, linked to the high-level needs of esteem and self-actualization. Both are achieved by engaging in activities such as learning, experimentation, research, but also by communicating with others in a social network.

One twist in comparing the emergent drives to Maslow's categories of human needs is that some needs, instead of serving one of the goals, may act as substitutes for our local control when it is unattainable, for example when there is uncertainty which we cannot eliminate. This includes the *social belonging* needs of bonding and emotional support, served e.g. by maintaining close kinship ties (Growiec and

Growiec, 2014). Accordingly, holding religious or superstitious beliefs is our paradoxical way to get around fundamental uncertainties and satisfy some of our *safety* needs without actually increasing local control.

3.4 Human Local Control in the Digital Era

Human control can be exercised over the natural environment and other humans. Over the hunter-gatherer, agricultural, and industrial era people have gradually learned to control the natural environment, but at the same we have dramatically intensified our contacts with other people. The latter fact mirrors both our huge reproductive success and the tremendous developments in communication technologies, which in the current digital era are allowing for instant information transmission and audiovisual contact with anyone across the globe. The consequence is that the digital-era human local control maximization process pertains mostly to interactions with other people and man-made machines.

One of the implications of this phenomenon is the democratization of access to information creation and sharing. Instead of relying on traditional ways of social status and hierarchy building, many of us now take the matters in own hands, managing extensive digital social networks and creating and sharing original media content on platforms such as YouTube, Facebook, Twitter, Instagram, Snapchat, etc. We are witnessing an unprecedented explosion of human creativity. Increased information creation and sharing brings measurable increases in aggregate human welfare: more useful data are created, collected and shared, more digital goods are consumed, and diffusion of useful technological knowledge is faster. And while some may dismiss many of those actions as economically unjustified (no direct monetary compensation), or view them as an epidemic of narcissism (Twenge and Campbell, 2009), or compulsive behaviors, the concept of local control maximization can actually rationalize these actions. I will elaborate more on this issue in Section 5.

4 Accumulation of Knowledge

In this section I will discuss how the human local control maximization process has led to systematic, cumulative technological progress. As we sought to expand the set of open choices and insure and diversify against risks, we naturally embarked on a path of knowledge accumulation. More precisely, knowledge accumulation was fueled by our emergent creativity (cognitive enhancement) drive.

But why have we succeeded at that? As argued above, evidence suggests that the key reason is the design of our brains. Namely, the capacity of our frontal cortex is sufficiently large for us to acquire the theory of mind and use it to create, document and share stories, gossip, tales, abstract ideas, and ultimately knowledge and technology (Tomasello, 2014; Harari, 2014). By pooling the collective memory among a wider group of people, the humankind managed to pass the threshold of systematic accumulation of knowledge. This capability translated into technological progress and ultimately economic growth because ideas are non-rivalrous, i.e. can be applied many times at once (Romer, 1986, 1990).

4.1 The Knowledge Accumulation Equation

To organize my narrative about the specifics of knowledge accumulation in the respective eras, I will use the following knowledge accumulation equation (an extension of the "idea production functions" used by Jones, 1999; Ha and Howitt, 2007; Madsen, 2008; Kruse-Andersen, 2017):

$$\dot{A} = \gamma F(K_A, L_A)A^{\phi} - \delta A, \qquad \gamma > 0,\ \delta \geq 0,\ \phi \leq 1. \tag{1}$$

I will focus on three parameters of the above equation: (i) the knowledge depreciation rate δ, (ii) the efficiency of knowledge creation γ, and (iii) the scale of operations, measured by the ratios of K_A and L_A relative to the respective world totals. To my knowledge, this focus is new to the literature. In contrast, I will be relatively less concerned with the magnitude of external returns to knowledge accumulation ϕ because it remains empirically unsettled – we do not even know if it is positive (representing "standing on shoulders" effects) or negative ("fishing out ideas"), Ha and Howitt (2007); Bloom, Jones, Van Reenen, and Webb (2017) – and crucially because my main message here is valid for any $\phi \leq 1$.

Please also note that the above specification assumes two types of inputs to the R&D (knowledge creation) process: R&D labor L_A, encompassing all the skilled work done by scientists and technical personnel, and *R&D capital* K_A. The latter factor, although typically disregarded in the R&D-based economic growth literature, makes a huge difference when comparing "idea production functions" across the four eras. The practicality and complexity of research equipment has undergone systematic, cumulative changes. The difference in usefulness of Ptolemy's astrolabe, Galileo's telescope, and the modern Very Large Telescope (VLT) is breathtaking; perhaps even more so is to think how early statisticians actually computed correlations and ran regressions over large datasets without relying on computers in their calculations.

I assume that F is increasing and concave in both factors, K_A and L_A. This implies scale effects in R&D (Kremer, 1993; Jones, 1995): the larger the economy, the faster the rate of knowledge accumulation, at least in the short run. F should be understood as an idea production function which is active within a certain technological paradigm (i.e., for a given γ). Observable technological progress then comes from incremental innovations which, in turn, rely on radical innovations for new research avenues to be opened (Olsson, 2000, 2005; Growiec and Schumacher, 2013). Even if there are "fishing out effects" within each technological paradigm (e.g., Bloom, Jones, Van Reenen, and Webb, 2017, suggest that $\phi < 0$), opening new paradigms rejuvenates technological opportunity (so that γ goes up) and may sometimes even begin a new era.

In the hunter-gatherer era, there was no conscious "research" activity, and thus discoveries could only be obtained via learning by doing, chance experiments, and informal thinking. Folk wisdom was transfered orally across generations, often in religious or superstitious frames, and generally was not shared beyond the tribe. As knowledge was stored only in people's memory, it was often forgotten. Thus the hunter-gatherer era knowledge accumulation process was characterized by (i) a high depreciation rate δ, (ii) low efficiency of knowledge creation γ, and (iii) small scale of operations. Though, compared to other animals, the *homo sapiens* already passed the threshold of a steady state, i.e. long-run stagnation in knowledge. As inferred from the excavations, the humankind's stock of knowledge began to be generally trending up already after the cognitive (Upper Paleolithic) revolution.

However, the further millennia brought dramatic improvements relative to this low starting point. These improvements affected all three key parameters of the knowledge accumulation equation.

First, a few key discoveries dramatically increased the durability of knowledge by reproducing it on "external memory", thus facilitating its storage and reducing the depreciation rate δ. These included the following: alphabet, writing and the printing press (invented in the agricultural era), telecommunications and audiovisual storage of data (in the industrial era), and ultimately digital memory and the Internet (in the digital era). Although writing, especially if exercised on highly durable material such as stone, theoretically made knowledge eternal, useful knowledge was still sometimes forgotten and had to be reinvented from scratch, as exemplified by the case of Roman bridges and aqueducts. Nevertheless, it may be argued that at least from the industrial era onwards the knowledge depreciation rate is essentially zero, as assumed by most of R&D-based economic growth theory (e.g. Romer, 1990; Jones, 1999; Ha and Howitt, 2007; Kruse-Andersen, 2017).

Second, the history of human knowledge witnessed also a number of breakthrough ideas which facilitated further discoveries and inventions, and hence increased the efficiency of knowledge creation γ. Examples of such breakthrough ideas are ancient philosophy, the university, the modern scientific method (empiricism, experimentation, mathematization, hypothesis testing), the research laboratory, industrial R&D, automation of tedious research tasks, digitalization of scientific knowledge, and the World Wide Web equipped with efficient search engines. All these developments gradually improved our ability to distinguish facts from myths, organize our knowledge set, locate gaps in it, ask novel, meaningful research questions, formulate testable theories, collect relevant data, and pursue empirical verification of our theories.

Furthermore, the efficiency of knowledge creation γ could also be affected by radical innovations which open up new research avenues. Focusing on last centuries, Gordon (2016) names three radical innovations which could have affected γ: (i) the inventions of steam engine, railroads, and cotton spinning (1750-1830) at the dawn of the industrial era, (ii) electricity, the internal combustion engine, and running water with indoor plumbing (1870-1900), and (iii) ICT technologies: computers, Internet, and mobile phones (since 1960), which initiated the digital era. The economic effects of the first two radical innovations lasted about a hundred years but – according to Gordon – the third one is probably going to be more short-lived. The key reasons, as I would speculate, are that nowadays international technology diffusion is much faster (Comin and Hobijn, 2010) and that the digital era is characterized by way higher growth rates (and shorter doubling times) than the industrial era and thus needs fewer years to completely reshape the world. As Gillings, Hilbert, and Kemp (2016) put it: "After RNA genomes were replaced with DNA, it then took a billion years for eukaryotes to appear, and roughly another two billion for multicellular organisms with a nervous system. It then took another 500 million years to develop neural systems capable of forming languages. From there, it took only 100,000 years to develop written language, and a further 4,500 years before the invention of printing presses capable of rapid replication of this written information. The digitalization of the entire stockpile of technologically-mediated information has taken less than 30 years. Less than one percent of information was in digital format in the mid-1980s, growing to more than 99% today".

4.2 Scale of Operations

The third parameter of the knowledge accumulation equation, the scale of operations, recorded perhaps a biggest increase in the course of human history. Its growth may have well been the foremost reason for the formidable acceleration in technological progress across the eras.

What I mean by *scale of operations* here is the degree of pooling of information, knowledge, and talent in the R&D process. Even if some researchers may be working alone, knowledge accumulation is always a cooperative process that occurs in a social network. And as knowledge is cumulative – previous knowledge acts as basis for new ideas – the size of this network affects the pace of technological progress. The most adequate empirical measure of the scale operations in global R&D is probably the size of the largest connected component of the global research network.

The network perspective is useful here because it helps distinguish between population growth, or growth in the number of active researchers, and connectedness of global social networks. Both elements grew over the millennia but started from a very low base. In the hunter-gatherer era, knowledge accumulation was pursued by small, disconnected bands and tribes, and hence there was massive duplication: the same knowledge (e.g., making fire, its use for cooking, the wheel) was acquired independently in different hunter-gatherer tribes across the world, and this knowledge was hardly shared. Low connectedness of social networks in prehistorical times explains also why agricultural revolutions happened independently in various parts of the world, separated by millennia over which the useful information and technologies (e.g., domesticated species) did not diffuse beyond surprisingly limited geographical areas (Diamond, 1997).

The game was changed at the time of great geographical discoveries and European colonial conquests, which began in the late 15th century. Europeans' ruthless "guns, germs and steel" (Diamond, 1997) had also the side effect of gradually connecting the previously fragmented social networks, unifying the global pool of technological knowledge, and increasing the scale of operations in the now global R&D process. This reduced duplication externalities in knowledge creation (Mokyr, 2002). The concentration of power in the hands of Europeans also facilitated the emergence of a universal language of science. Moving away from Latin, it became a mixture of French, German, and English, until it was fully dominated by English in the second half of the 20th century. Yet another unifying factor was the broad adoption of common mathematical methods of scientific research.

I should also remark that the scale of operations in global R&D grew (and is still growing) also because the share of population and capital attributed to R&D actions is arguably increasing over time at the global scale. Given the digital-era emphasis on R&D, technological startups, and innovation-led growth, this trend is likely to persist and even further intensify in the future.

4.3 R&D in the Digital Era

Knowledge accumulation rapidly accelerated in the digital era compared to the industrial one. Taking the local control maximization perspective, one reason for this acceleration is that the digital era puts *useful data* on center stage, and therefore obtaining new knowledge becomes often a goal in itself, instead of an intermediate goal en route to increasing *value added*, the key variable of the industrial era, or increasing *agricultural production*. This encourages the world economy to allocate more resources to R&D than ever before.

The second reason is the increased efficiency of knowledge creation. Radical innovations such as computers, Internet, and mobile telephony, opened up new exciting research avenues. At the same time, the digital era also produced new, highly specialized forms of R&D capital. Many inventions would have never been obtained if not for sophisticated, computerized physical, chemical and pharmaceutical laboratories and the abundance of general-purpose computing power. In the digital era, computational complexity is less and less an issue, allowing us to build more sophisticated models of reality, to pursue more detailed empirical identification of facts based on statistical data, and to engage in extensive exploratory research and data mining. Digital-era technologies also allow us to instantly search the massive base of earlier publications.[1]

Extrapolating past trends into the future, we should then expect rapid knowledge accumulation with an ever larger fraction of inputs assigned to the R&D process, and an increasing contribution of R&D capital to the newly created knowledge. Several authors (e.g., Frey and Osborne, 2013; Brynjolfsson and McAfee, 2014; Acemoglu and Restrepo, 2016) suggest, however, that this may not be the end of the story. Instead, we may be now just in the advent of an explosion of a new "game-changing" radical innovation in the field of artificial intelligence (AI). Sufficiently advanced AI may soon enter the "idea production function" not merely as R&D capital, but may in fact substitute out humans in (at least part of) their researchers' tasks.

Brynjolfsson and McAfee (2014) predict that sophisticated AI technologies will likely turn out decisive for growth dynamics in the near future by developing "gradually, then suddenly", fueled by their highly scalable character and – potentially – ability to self-improve. Hence, even though AI is still nascent at the moment,[2] it is easy to imagine that highly developed machine learning and big data algorithms, autonomous laboratories, automatic translators, text generators, and multi-function robots, can potentially have a massive impact on knowledge accumulation in the future.

This perspective also raises the question of substitutability between computer software (which includes AI algorithms) and human R&D labor. Thus far they have been complementary: software was a tool in researchers' hands. The primary reason is that so far computer algorithms have been very bad at *ideation*, creativity, or asking useful research questions (Brynjolfsson and McAfee, 2014). But this does not have to be always the case. Day by day, AI algorithms are getting better and better at pattern recognition based on big data, classification, categorization of various sorts of content, and making adaptive decisions in noisy, variable environments – and they are much faster than humans at all that.

The role of AI in future R&D processes depends on the answers to two following questions. First, is ideation a sophisticated incarnation of pattern recognition or a qualitatively different feature? Some preliminary results suggest the former answer: AI algorithms are already able to write fiction books, compose music or draw

[1] The quantification of the role of R&D capital for R&D output and economic growth is still a gap in the literature, though. I suppose that this is the case because R&D capital emerged as a significant contributor to R&D output relatively recently.

[2] Although AI is already able to systematically outperform *all* humans in such sophisticated games like chess and Go (DeepMind AlphaZero), and *Jeopardy!* (IBM Watson), safely and efficiently drive cars in regular traffic (e.g., Google car or Tesla), and support physicians in diagnosis (again IBM Watson).

artistic pictures (Schmidhuber, 2009a). This would imply, however, that there is no qualitative difference between human cognitive abilities and computer software in R&D. If this is true, sufficiently developed AI algorithms given sufficient computing power may one day surpass humans in doing research.

The second question pertains to the expected pace of AI progress. Namely, how high are the *returns to cognitive reinvestment* in machine intelligence? (Yudkowsky, 2013). How efficient will the future AI be in re-designing itself and its environment in order to improve its research skills? Humans are in this regard limited by their innate cognitive capacity. We are unable to rewire our brains, and so we circumvent this limitation by increasingly relying on external memory, data collection equipment, and computational power. We also increasingly pool our resources by working in ever larger research teams whose members have increasingly specialized sets of skills. As our knowledge set is growing but our brains are not, interdisciplinary "Renaissance Men" are long gone (Jones, 2009). Unfortunately, speed and accuracy of our interpersonal communication are far from perfect, and thus we may be missing plenty of interdisciplinary insights. AI algorithms running on fast computers, in contrast, communicate extremely fast and without error. They also by far surpass us in terms of speed and serial depth of computation (Hanson and Yudkowsky, 2013). In contrast to us humans who cannot rewire our brains, machine intelligence is also (at least theoretically) potentially able to recursively rewrite its code provided that it is able to prove that the rewrite is beneficial (Schmidhuber, 2009b). So far, AI is however markedly lagging in terms of versatility and adaptivity. If this is resolved, we may observe a rapid buildup of AI skills, and even an intelligence explosion. Whether it is going to be good or bad for us, I will try to elaborate in Section 7.

5 Measuring the Level of Development

So far I have argued that systematic progress in human local control maximization manifests itself as knowledge accumulation. In this section I will describe how it is transformed into economic growth. I will review the key differences between the eras and argue that the level of global economic development ought to be measured differently in the respective eras. This postulate has intriguing implications for the ongoing digital era.

5.1 Level of Development Across the Eras

How should one measure economic development at the regional, sectoral, national, and global level? Economists' first choice would be to use the GDP, gross value added (GVA), or some related concept from the National Accounts. Moreover, when dealing with pre-industrial eras, we would typically voice discontent about the lack of exact data, and proceed to construct proxy variables based on the available information on population size, prices of traded and non-traded goods, etc., and use some creative interpolations. This is precisely what is being successfully accomplished by, e.g., the Maddison Project.[3] I have absolutely no objections to this approach; yet, taking a step back and looking at the whole human history suggests that it may provide an incomplete characterization of economic development at the global scale and in the very long run. From this perspective, it appears to be a struggle to provide *exact* measurement of only *part* of our actual development.

If one accepts the proposition that global economic development is a consequence of the human drive to maximize local control, a direct way to measure development at the global scale would be to calculate the grand *sum of human control*. But what could this mean in practice? We need something that is easier to handle empirically. Proxy variables like global population size, global GDP, multidimensional measures of the degree of satisfaction of certain kinds of needs, subjectively reported well-being, aggregate optimization power of human minds and man-made machines, the aggregate stock of non-redundant data, etc., all have problems of one kind or another.

My proposition is therefore to use era-specific measures of aggregate human control. With the wisdom of hindsight, I argue that in the hunter-gatherer era the extent of human control could be identified with the total carrying capacity of ecosystems under human rule. I call it the habitat capacity, Hab. In the agricultural era, characterized by Malthusian population dynamics, the extent of human control was in turn equivalent to total human population, POP. In the industrial era, the extent of human control could be identified with world GDP. For the digital era I do not have the wisdom of hindsight, so I cannot tell for sure. Nevertheless in the following paragraphs I would like to propose to equate digital-era aggregate development Q with flows of *bits of useful data*. To my knowledge, this is a new perspective.

Aggregate development Q can then be usefully decomposed as

$$Q = Hab \cdot \frac{POP}{Hab} \cdot \frac{GDP}{POP} \cdot \frac{Q}{GDP}. \tag{2}$$

Equation (2) signifies that each technological revolution opens up a new dimension of development, which was previously fixed. Moreover, growth in the new dimension *adds* to the growth in the earlier dimensions (and does not substitute them). In fact, thanks to positive feedback effects each new era strongly accelerates previous-era growth.

The logic of this decomposition requires that a new technological revolution can potentially occur only if the previous-era development level is sufficiently high. The

[3] https://www.rug.nl/ggdc/historicaldevelopment/maddison/

new era always operates *on top of* the previous one: there would have been no agricultural revolution if the habitats had not been conquered by the earlier hunter-gatherers; no industrial revolution without secure supply of food for a large (and growing) population; and no digital revolution without reliable supply of both food and a variety of industrial-era goods and services. The consecutive revolutions can be also viewed through the lens of Maslow's (1954) hierarchy of needs: each era opens new possibilities for satisfying higher-level needs, but pursuing them is possible only if lower-level needs are already sufficiently satisfied.

In the hunter-gatherer era, successful development could only be done by acquiring control over new habitats which were previously occupied by other species. Humans gradually eliminated their enemies and advanced up the food chain (Harari, 2014). Once able to do this, they gradually spread around the world. The amount of food and other natural resources which they could draw from a given piece of land were roughly fixed because their technology did not allow them to systematically, intentionally transform ecosystems in their favor. Therefore population size given habitat capacity was roughly fixed, and so was GDP per capita (at subsistence level). Hence $Q \approx Hab$.

The agricultural era, in contrast, was founded on the idea of transforming habitats. Domestication of food crops opened up the intensive margin of land use, allowing to sustain more human lives per square mile. Successful development could then be done not just by conquering new land, but also by turning already acquired land into farmlands and intensifying agriculture (Diamond, 1997). On the other hand, the agricultural era was a Malthusian epoch: all increases in output were eaten up by increases in population size, and thus GDP per capita was roughly fixed (at subsistence level). Hence $Q \approx POP$.

The technological breakthroughs of the industrial revolution brought a new form of economic development, via industrial production and accumulating physical and human capital. This opened up the intensive margin of consumption per capita, allowing more goods (in terms of quantity, quality and variety) to be consumed per person. After the initial period of rapid capital buildup, industrial-era economies began to exhibit an ever growing demand for skilled labor. This opened a children quantity–quality trade-off and triggered a decline in fertility, thus breaking the Malthusian mechanism. In sum, in the industrial era aggregate human control could be increased either extensively (via population growth) or intensively (via growth in GDP per capita). On the other hand, information content per unit of GDP was roughly fixed in the industrial era because data storage was limited to analog means, primarily paper, and later audio and video tapes (Hilbert and López, 2011). Hence $Q \approx GDP$.

The key novelty of the digital-era technology is that it advances our local control by facilitating collection, transformation, and communication of data. Compared to the industrial era, this opens up a new intensive margin of accumulating *bits of useful data* per unit of GDP. It has never been so easy to obtain useful information than in the digital era, and we are systematically taking advantage of this e.g. by spending many hours per day browsing the Internet. At the same time, we are increasingly active as creators of online content. It is getting clear that we tend to draw more utility from useful data than its pecuniary value would suggest. In sum, I posit that in the digital era aggregate human control can be increased either extensively (via GDP growth) or intensively (via growth in useful data per unit of GDP). Hence Q is roughly equal to the total flow of bits of useful data.

At this point, I would like to emphasize that there is no theoretical reason to believe that the accumulation of useful data *is* human local control. In fact, saying so would amount to adopting *dataism* which "declares that the universe consists of data flows, and the value of any phenomenon or entity is determined by its contribution to data processing" (Harari, 2017, p. 428). In contrast, I posit that useful data are not the ends, but means to obtaining more human control. I am open to the possibility that another technological revolution would come in the future, opening yet another dimension of economic development and justifying a split of Q into further factors. There are already a few possible candidates for such breakthrough technologies: artificial general intelligence (AGI), nanotechnology, quantum computing, etc.

5.2 Feedback Loops Between the Eras

What blurs the sharp distinction between the consecutive eras is that they co-exist in time and are interrelated in data. Each new era opens an entirely new dimension of action but also applies the new developments and new knowledge to the previous-era economy. Such positive feedback effects were visible across all eras.

Firstly, the agricultural revolution allowed our ancestors to transform ecosystems in order to make them more habitable. Several previously hostile and thus sparsely inhabited biomes such as steppes may have suddenly become fertile once domesticated crops and animals have been introduced. Therefore the hunter-gatherer measure of economic development, habitat capacity Hab, was increasing in the agricultural era (until the point of satiation).

Secondly, the industrial era not only opened up a new sector of the economy, but also revolutionized the old one: farming. Mechanization of agriculture, nitrogen, phosphate and potassium fertilizers, and the Green Revolution contributed to massive increases in crop yields. At the same time the industrial-era progress in medicine implied a massive drop in mortality. Both effects combined led to a huge improvement in the agricultural era measure of success, i.e., total population (and population density). However later fertility dropped, and nowadays rapid population growth is observed only in poorest regions where people are still employed in subsistence agriculture.

Thirdly, modern digital-era technologies massively improve the efficiency of industrial production. There are many channels: cost-reducing and quality-enhancing automation, increasing quality and variety of goods thanks to the featuring of electronics, improved information flow within and among firms, opening of new service sectors (business-to-business, ICT services, etc.), etc. Furthermore, improved information flow facilitated rapid globalization, global fragmentation of production and the creation of global value chains, all of which were also efficiency-improving innovations. Robert Solow may have said in 1987: "You can see the computer age everywhere but in the productivity statistics" (Solow, 1987), but nobody would dare to say this today. Owing to a positive feedback loop, the effects of the digital era have become very visible in productivity statistics.

5.3 Documenting Subsequent Accelerations and "Secular Stagnations" Across the Eras

The pace of economic development has been strongly accelerating across the eras. The doubling period for aggregate human control Q has shrunk from 34 360 years in the hunter-gatherer era to 885 years in the agricultural era (population growth), 33 years in the industrial era (total GDP growth in PPP), and 2.8 years in the digital era (growth in the volume of communicated data). What fueled these impressive accelerations? As I explained above, in my view the two key factors were the improvements in the knowledge accumulation process and the sequential opening of entirely new dimensions of economic development. Knowledge accumulation was accelerated by radical innovations which reduced the rate of knowledge depreciation, increased the returns to research effort, accelerated diffusion of information, and – most importantly – increased the scale of operations in R&D, measured by the largest connected component of worldwide research network. Opening new dimensions of economic activity, in turn, was associated with opening new technological paradigms, driving up the pace of knowledge accumulation, at least until it got overburdened by "fishing-out" effects (Jones, 2005; Bloom, Jones, Van Reenen, and Webb, 2017). Upon each technological revolution, the new economy grew beside its previous-era counterpart, and there were also sizable positive feedback loops between eras.

The flipside of the aforementioned accelerations were "secular stagnations" which appeared in the previous-era economy whenever the positive feedback loop from the next era had subsided, and further development remained concentrated in the last era. The topic is hotly discussed in the context of GDP growth slowdowns currently observed across the developed economies (Gordon, 2016). In my view, the "secular stagnation" hypothesis can be intuitively explained in the current framework as hypothesized arrival of a period when positive feedback effects from the digital-era economy on the industrial economy have subsided and further growth remains concentrated only in the digital era.

Economic growth in the hunter-gatherer era was driven by two factors: acquisition of new natural habitats, coupled with proportional increases in human population, and gradual accumulation of technological knowledge. Both factors grew at a snail's pace. According to Kremer (1993) data, the population growth rate between 25 000 BCE and 5 000 BCE amounted to 0.002% per annum. As the resultant population doubling period amounted to 34 360 years, the hunter-gatherer population failed to double in this period.

In the agricultural era, in contrast, new technological developments allowed to transform natural habitats, intensify food production, and hence increase population density. At the same time, knowledge accumulation accelerated thanks to a few breakthrough ideas, like writing or philosophy, and to the gradually increasing scale of operations in R&D. There was also a positive feedback loop with the hunter-gatherer economy in the sense of acquiring additional habitable land. However, as eventually all useful land was subjected to human control, a "secular stagnation" in populated land area set in. According to Kremer (1993) and Piketty (2014) data, in the period between 5 000 BCE and 1820 CE population growth rate amounted to 0.078% per annum (implying a doubling period of 885 years). While still modest by modern standards, this amounted to a massive 39-fold increase in the population growth rate compared to the hunter-gatherer era. Owing to the Malthusian mechanism, though, increases in human control per capita Q/POP were very limited. According to Piketty (2014) data, GDP per capita growth rate (PPP) in 0–1820 CE amounted to just 0.019% per annum (so that it would potentially double only in 3600 years).

The industrial revolution brought a further acceleration in economic growth, through a number of channels. First of all, the world was already connected at the time, so knowledge accumulation could operate at a very large scale. Second, after the printing press knowledge depreciation was essentially driven down to zero. Third, breakthrough industrial technologies set off the processes of physical and human capital accumulation, as well as increases in GDP per capita and (widely shared) standards of living. There was also a positive feedback loop with the agricultural economy: mechanization of agriculture and spread of fertilizers increased crop yields, while developments in medicine, hygiene and increases in standards of living dramatically reduced mortality, paving the way for a demographic explosion (Boucekkine, de la Croix, and Licandro, 2003). Soon thereafter, however, the children quantity–quality trade-off associated with rising educational attainment triggered a decline in fertility, gradually bringing population growth rates down, eventually to zero or even mildly negative values in the Western world, a clear "secular stagnation" in population growth.[4] According to Piketty (2014) data, in the period 1820-1990 CE the headline industrial-era measure of development, world GDP (in PPP), grew at a rate of 2.14% per annum, doubling every 33 years. The intensive margin of growth, GDP per capita, grew at a rate of 1.17% per annum, doubling every 60 years. At the same time, global population grew at a rate of 0.96% per annum, doubling every 72 years.

The acceleration brought by the digital revolution, again, worked through two channels: accelerated knowledge accumulation and the opening of a new dimension of economic growth. Knowledge accumulation was accelerated thanks to the development of digital R&D equipment, facilitating knowledge storage and retrieval, connecting the researchers via the Internet, and automating the most tedious research tasks. In terms of the new dimension of growth, in turn, the digital revolution set off the processes of rapid hardware and software development. Furthermore, a strong feedback loop was activated with the industrial era via automation, increasing quality and variety of goods, improved information flow, and the opening of new sectors. According to Hilbert and López (2011) data, between 1986 and 2007 processor capacity grew at 58% per annum (doubling every 1.5 years, the exact Moore's Law), data storage grew at 23% per annum (doubling every 3.3 years), and data communication grew at 28% per annum (doubling every 2.8 years). Owing to the positive feedback loop with the industrial-era economy, GDP growth seems to have accelerated in the digital era. According to Piketty (2014) data, in 1990-2012 CE the average GDP per capita (PPP) growth rate was equal to 2.08% (doubling every 34 years).

This is where we are now. Are we going to witness a gradual decline in the feedback effects from the digital era to the industrial era in the future? Hard to tell, but if we were to extrapolate from the previous technological revolutions, the answer would be positive: yes, there would be a "secular stagnation" in GDP growth. Nevertheless we would still be observing soaring growth rates of information creation, storage, and communication – and ultimately, aggregate human control would be growing faster than ever. Until we perhaps reach another technological revolution.

Thinking of possible further technological revolutions, one should bear in mind that although exponential growth rates and doubling times are a convenient descriptive tool, exponential growth should not be expected to continue indefinitely. That would call for making at least one highly unlikely knife-edge assumption (Growiec, 2007). Also factually, looking at the millennia of past economic growth suggests that sigmoid (S-shaped) long-run development patterns are in fact much more accu-

[4]However, according to Piketty (2014) data, in 1990-2012 population grew at an unprecedented rate of 1.3% per annum, implying a doubling period of 53 years. This aggregate figure was driven exclusively by population growth in relatively backward, largely agricultural societies of e.g. Sub-Saharan Africa or Southeast Asia.

rate descriptions of reality. Habitat acquisition has been completed long ago, global population growth is projected to stop in the next century, and it is conceivable that global GDP will eventually stagnate, too. What is striking, though, is that the sigmoid functions describing each consecutive era are getting sharper, implying faster growth soon after inception of the era, and that the consecutive revolutions are getting closer to each other in time. Extrapolating this observation forward in time, Hanson (2000) speculates that "within the next century a new mode might appear with a doubling time measured in days, not years."

Moreover, using an alternative methodology of estimating a super-exponential growth pattern with power-law acceleration in the growth rate, Johansen and Sornette (2001) find that the data are consistent with a spontaneous singularity in the year 2052 ± 10, signaling an abrupt transition to a new regime. Similar super-exponential growth patterns in ICT data have been documented by Nagy, Farmer, Trancik, and Gonzales (2011).

The *singularity* (Kurzweil, 2005) is however an elusive concept. The literature abounds with futuristic interpretations of singularity, going way beyond the formal mathematical definition of a vertical asymptote, i.e. the situation where a certain variable (e.g., GDP per capita, or aggregate development Q) tends to infinity in finite time. This seeming infinity is frequently identified with qualitative change but it may also be an artifact of poor approximation of sharply rising but finite-valued time paths with hyperbolic curves which possess a vertical asymptote.

5.4 Designing New Measurement Methods for the Digital Era

If we agree that the digital-era economy is able to increase the extent of human control (and the degree of satisfaction of human needs) beyond the pecuniary value of its goods and services, it must follow that our standard measures of economic development, such as GDP, are adequate for the industrial era but obsolete when dealing with the digital era.

The concept of the National Accounts is rooted in the industrial era. This reflects both its history and its focus on the *dollar value* of inputs and outputs of production processes. By construction, differences in information content per unit of value added are not acknowledged in the National Accounts. There are reasons to believe that these differences carry an economic meaning, though. For instance, several classes of digital goods, such as online entertainment (free music, movies, etc.), open source and freeware utility software, or open access scholarly publishing, are entirely missed by the National Accounts even though they contribute to the well-being of the society (Brynjolfsson and McAfee, 2014), satisfaction of human needs, and ultimately the aggregate span of human control.

Therefore, wishing to account for the tremendous growth in the world's cumulative ability to compute, store, and communicate useful information, by far exceeding GDP growth rates in the last decades, here I tentatively propose that in the digital era, economic development Q should be associated with flows of *bits of useful data* among humans and devices. This communication can be done through traditional means (e.g., talking, writing) but also increasingly through digital means. I think that insisting on looking at global economic development through the lens of GDP growth misses the fact that in the digital revolution – for the first time in history – decoupled total data transmission from GDP.

The tremendous growth of the digital-era economy justifies the need to design a counterpart of National Accounts specifically for data flows instead of flows of goods and services. To account for digital output, we need a concept such as *bits of useful data* which would work like value added in the National Accounts. In particular, flows of *useful* data should be distinguished from all data flows, perhaps in a manner similar to the distinction between value added and global output. One should also redefine the sectoral breakdown of the economy in order to get a more precise measurement of the structure of data flows. This will clearly be a huge endeavor, but likely not an impossible one, and clearly a very useful one. Beyond economics, preliminary ideas on how this could be accomplished have been proposed by information theorists (Hilbert and López, 2011; Hilbert, 2017).

This call for a new measurement design comes with a number of caveats. First, data flows rarely respect national boundaries, thus suggesting that the measurement ought to be done at the global scale. Therefore this is in fact a call for Global (not National) Data Accounts. Second, there is a lot of redundancy in data transmission, storage, and even creation. A lot of thought must be devoted to the filtering of raw data flows so as to distill the key variables in question. Third, some data flows may be unintended or detrimental to the users. For one thing, think of the spread of fake news and Internet hate. I will elaborate more on this in Section 7. Fourth, the data accounts would have to cleverly embrace the evolving role of intellectual property, and work its way around data secrecy in order to adequately compute the volume of data flows without compromising the proprietary character of some datasets.

Of course, this is not to say that National Accounts should be discontinued. To the contrary, in fact the creation of Global Data Accounts may *improve* the accuracy of National Accounts. I am only saying that the digital revolution opened up a new dimension of actions which has not been sufficiently measured yet. So far we can indirectly evaluate the digital-era economy by accounting for the feedback effects of the digital era on the industrial era, which we do by using data on GDP and value added (e.g., Jorgenson and Stiroh, 2000; Timmer and van Ark, 2005). We fail at capturing the total effects of digital development, though.

It is also critical to distinguish between inputs and output of the digital-era production process. Flows of useful data are generated using the previously accumulated stocks of hardware and software, both of which can and should be measured. I will elaborate more on digital-era production inputs in Section 6.

5.5 Preliminary Evidence on the Data Explosion

While probably insufficient for economic purposes, information theorists have already produced their first assessments of the level and growth of the world's total storage, transmission and computation of data. The seminal contribution of Hilbert and López (2011) tracks 60 analog and digital technologies during 1986–2007 and announces that: "In 2007, humankind was able to store 2.9×10^{20} optimally compressed bytes, communicate almost 2×10^{21} bytes, and carry out 6.4×10^{18} instructions per second on general-purpose computers. General-purpose computing capacity grew at an annual rate of 58%. The world's capacity for bidirectional telecommunication grew at 28% per year, closely followed by the increase in globally stored information (23%)."

To put these numbers in perspective, digital storage, transmission and computation of data has also been compared with its analog counterpart, done for millennia in human brains. In this regard, note that before the invention of writing, human brains were the only form of store of data. After the radical innovations of writing, printing press, analog means of audio and video recording, etc., data storage was gradually outsourced to "external memory", and the stock of stored data per person began to rise. The digital revolution put this growth to an entirely different gear, though. Gillings, Hilbert, and Kemp (2016) write: "Evolution has transformed life through key innovations in information storage and replication, including RNA, DNA, multicellularity, and culture and language. We argue that the carbon-based biosphere has generated a cognitive system (humans) capable of creating technology that will result in a comparable evolutionary transition. Digital information has reached a similar magnitude to information in the biosphere." Most of this explosion took place in the last three decades as the world's capacity to store, communicate, and compute information has begun to soar only since the 1980s (Hilbert and López, 2011). Quantitatively, "[i]nformation technology has vastly exceeded the cognitive capacity of any single human being (...). In terms of capacity, there are two measures of importance, the number of operations a system can perform, and the amount of information that can be stored. The number of synaptic operations per second in a human brain has been estimated to lie between 10^{15} and 10^{17} (...). While this number is impressive, even in 2007, humanity's general purpose computers were capable of performing well over 1×10^{18} instructions per second (...). Estimates suggest that the storage capacity of an individual human brain is about 10^{12} bytes (...). On a per capita basis, this is matched by current digital storage (5×10^{21} bytes per 7.2×10^9 people)." (Gillings, Hilbert, and Kemp, 2016)

A few observations are due here. First, these estimates, while impressive, are exact only up to an order of magnitude. This is not yet the statistical precision we are accustomed to in economics.

Second, these estimates include a mixture of input and output, stock and flow variables. Data storage capacity and processor computational capacity can be thought of as stock input variables, akin to physical capital used in industrial production. The amount of actually communicated data is, on the other hand, a flow output variable. It is however not yet our postulated concept of *flows of useful data*. To reach that level of precision, one would have to filter out the data redundancy as well as intermediate, auxiliary data used up in computations. In economic terms, total communicated data is a digital-era equivalent of global production, whereas flows of useful data – an equivalent of value added.

Third, setting up Global Data Accounts and identifying the flows of useful data may also indirectly facilitate measurement of global knowledge A. Thus far the latter variable has proven to be inherently difficult to measure. The best proxies we have as economists are probably total factor productivity (TFP), patents, and research articles plus scientific books. They are all very crude measures of technological knowledge, though. TFP is a residual measure which also lumps all sorts of measurement and specification error. Direct measures like patents or research articles and books, in turn, are problematic because of incomplete treatment of actual usefulness and information content of each patent, article or book.

6 Accounting for the Key Factors of Production in the Respective Eras

Humans have learned to produce a variety of goods and services, using a variety of inputs. To organize this variety and enable simplified modeling, macroeconomists frequently group the inputs and outputs into relatively homogeneous clusters and plug them in an aggregate production function. Whether this concept is an adequate and useful representation of reality, remains open to debate (Felipe and Fisher, 2003; Temple, 2006). For one thing, aggregate production functions always feature some residual terms such as TFP (Solow, 1957). In this section, though, I use the aggregate production function representation of production processes in order only to organize the thinking rather than to produce exact implications.

6.1 Aggregate Production Function Across the Eras

Throughout the following analysis I describe the production process with an aggregate production function of form:

$$Q = F(A, K, L), \tag{3}$$

where Q is the output variable representing aggregate human control, A is technological knowledge, K is the "capital" input, and L is the "labor" input. In accordance with all literature I suppose that F is increasing in A, K, L. I also assume constant returns to scale with respect to K and L but not A (Romer, 1986, 1990; Barro and Sala-i-Martin, 2003). Across the four eras, the catch-all terms K and L will encompass very different inputs, described below. To my knowledge, thus far the literature has not systematically investigated the changes in composition of the key production factors across the entire human history.

I posit that K and L are gross complements, as capital and labor have been over the last century (Chirinko, 2008; Klump, McAdam, and Willman, 2007, 2012; Oberfield and Raval, 2015) and as land and labor naturally are in the agricultural production function. Gross complementarity means that an increase in the quantity of one factor implies an increase in the elasticity of output with respect to the other factor. In a perfect-competition setting, the latter also means an increase in the other factor's share of output. For example, accumulating more capital depresses the capital's share of output in favor of labor. Gross complementarity implies also that each of the factors is essential for production, i.e., that if its supply is zero, so is output.

The output variable of the hunter-gatherer era is total food production from habitats under human control. It is roughly equal to the natural habitat capacity, Hab. Flows of food and other natural resources (e.g., furs for clothing, wood for shelter, healing herbs, etc.) from a given piece of land are essentially fixed. Therefore the "capital" factor can be identified with cumulative carrying capacity of natural habitats under human control, $K = Hab$, and "labor" is proportional to population size which, in turn, is also proportional to the habitat capacity, so that $L = Hab$. Population size and labor hours per person are set from the binding subsistence constraint. Hence,

$$Q = F(A, Hab, Hab) = Hab \cdot \bar{F}(A). \tag{4}$$

In consequence, the slow increases in output Q per piece of land are exclusively due to increases in A, i.e. due to the invention of better stone tools, hunting methods, new uses of fire, etc.

The output variable of the agricultural era is total food production, and the inputs are land ($K = Land$) and agricultural labor ($L = POP$):

$$Q = F(A, Land, POP) = POP \cdot \tilde{F}\left(A, \frac{Land}{POP}\right). \tag{5}$$

Thanks to the domestication of a range of food crops, production of calories per land area was higher in the agricultural era than in the hunter-gatherer era, allowing to feed more people. In the early years of the Neolithic revolution farmland was relatively easy to accumulate, and thus people started gradually transforming natural ecosystems to arable land. This generated food surpluses, allowing to increase population size and accumulate more agricultural labor. These food surpluses also drove the rise of non-food-producing specialists, emergence of landed elite, and acquisition of more agricultural land. As the supply of yet uncultivated land began to shrink, and land became the limiting factor of production, growth in output per worker declined. Following the Malthusian mechanism (Kremer, 1993), consumption was driven down to subsistence levels (except for landed elite, and especially if this elite practiced primogeniture, Bertocchi (2006)). What followed was an era of few wealthy feudal lords and many poor peasants (who were often tied to their lords in serfdom). Scarcity of land implied high land-owner shares of agricultural output and great inequality (Piketty and Zucman, 2014). When the average cultivated land area per person stagnated, all further increases in output per person were driven by technological progress, captured by increases in A.

The output variable of the industrial era is GDP, and the inputs are physical capital K and human capital H:

$$Q = F(A, K, H). \tag{6}$$

In the early years of the industrial revolution, physical capital was relatively scarce and easy to accumulate. The first capitalists began to build factories, manufactures, assembly lines, and set up their enterprises. This increased output per worker, making labor the scarce input, at which point the demand for industrial labor went up. This demand rise applied in particular to the services of *skilled* labor. The supply-side response was increased educational attainment, partly due to the emergence of public education, ending up in secular increases in human capital per worker. The equilibrium response was an upward trend in wages and the emergence of skill premium. The capital share of output, and income inequality, first went up (in the Marxian period) and then down (Galor, 2005, 2011).

In mature industrial economies, most of the 19th and 20th century was characterized by Kaldor's (1961) "stylized facts", with constant factor shares of output, a steady return to capital, and an upward trend in wages, capital and consumption which grew at the pace of output. The driving force of long-run economic growth was technological progress captured by sustained growth in A, further accompanied by marked increases in physical and human capital per capita over a prolonged transition period (Barro and Sala-i-Martin, 2003; Jones, 2005). These "stylized facts" ceased to hold only around 1970s–1980s, immediately prior to the dawn of the digital era (Jones and Romer, 2010).

The industrial era also influenced the factors of production in the agricultural-era economy. The emergence of the early capitalists as well as mechanization of agriculture somewhat diluted the power of the incumbent landed elites, eventually reducing agricultural-era inequality and democratizing control of land, thus at least partially preparing ground for Revolutions of 1848.

The output variable of the digital era, as argued above, is the flow of useful data, and the inputs are hardware (in place of the "capital" input K), and software (in place of the "labor" input L):

$$Q = F(A, Hardware, Software). \tag{7}$$

In equation (7), hardware includes all sorts of digital capital goods (computers, cell phones, various electronic devices, industrial robots, etc.), whereas software includes all sorts of algorithms and programs which can be run on these devices. Crucially, the "software" amalgamate includes also human labor because it is human programmers who ultimately write the instructions to be performed by hardware. Human labor in the digital era is thus complementary to hardware and substitutable to (pre-programmed) software.

In the beginning of the digital era, hardware was relatively scarce and easy to accumulate, and therefore people started accumulating it, thus augmenting their skilled labor. This brought about fast increases in computation power (at the pace of Moore's Law) as well as flows of useful data, both per person and per unit of employed software. At the same time, thanks to scalability and complementarity with rapidly expanding hardware, leading software producers such as Microsoft, Google and Facebook have become global superpowers. Compared to the mature industrial era, three last decades also brought about gradual declines in the labor share of global output (Karabarbounis and Neiman, 2014), increases in the profit share and CEO pay (Gabaix and Landier, 2008; Barkai, 2017), and increasing inequality, in particular top income inequality (Piketty and Saez, 2003; Piketty and Zucman, 2014; Jones and Kim, 2017).

This is where we are now. But how will the digital-era economy develop in the future? Economic logic as well as analogies to the history of agricultural and industrial eras suggest that we may be soon observing rapid increases in the demand for (broadly defined) software. This is because software has by now clearly become the scarce factor, the bottleneck of the digital economy. With a few exceptions, our software generally does not use the available computing power efficiently. We have even coined a humorous saying on this: "What Intel giveth, Microsoft taketh away". In fact many fancy software features that we like, such as prettier graphics, more intuitive frontend solutions, etc., generally take away processor power and burden memory with information which is not central for the core digital process. Furthermore our software is very specialized and rigid in terms of data requirements. This means that, despite tremendous growth in available computational capacity, the level of automation of the digital-era production process remains low. Most of the time when computing power stays idle is when the user has to set the instructions manually.

Because computing power is accumulated at the pace of Moore's Law, the price of software relative to hardware is quickly going up. This creates potentially huge economic rewards to developing better software, in particular AI. I am pointing specifically at AI here because it appears the best way to circumvent the relatively slow pace of human thinking and computer–human communication. Both elements slow down hardware substantially compared to within-computer computation and computer-computer data flows. The largest increases in digital-era output are observed when tasks are fully automated.

Thus far there has been place for both humans and computer algorithms in the "software" amalgamate, but it does not have to always be the case. Algorithms are vastly faster and more accurate, but they lack the cross-domain versatility of the human mind (Yudkowsky, 2013) and the ability of ideation and creativity (Brynjolfsson and McAfee, 2014). But, as mentioned earlier, they are getting better and better at pattern recognition based on big data, classification, categorization of various sorts of content, etc. Automation is already gradually eliminating routine jobs, both manual and cognitive (Acemoglu and Autor, 2011; Autor and Dorn, 2013; Frey and Osborne, 2013), and the remaining jobs are safe only until the development of sufficiently sophisticated and versatile AI technologies. Muehlhauser and Salamon (2012) present the key advantages of AI relative to humans: increased computational resources, vastly superior communication speed, increased serial depth (of thought), duplicability, editability, goal coordination, and improved rationality. These advantages are important for gauging why AI is a potentially valuable substitute for human capital in performing majority, if not all tasks of the digital-era economy. Even so, people tend to protect themselves against full automation of human tasks by creating new jobs (Acemoglu and Restrepo, 2016). Indeed, nowadays people enter e.g. into creative niches of YouTube vlogging, Instagram modeling, videogame graphics art, etc., at a quite high rate. But will we be always able to do so? I will try to comment on this in Section 7.

I also expect that the internal structure of the "software" amalgamate will also shape the future changes in income inequality. So far the digital era has been marked by increasing income inequalities and profit shares. To some extent, this may be exploiting the fact that computer software is extremely highly concentrated nowadays (e.g., Microsoft, Google, Facebook), owing to the fact that software is highly scalable and its market is global. At the same time, it is the bottleneck of the digital era, relative to the more dispersed and often idle hardware. If computer software will indeed be developed in the future sufficiently strongly to contribute a larger share of the total "software" input, income inequality and profit shares will rise further. These trends will likely be even exacerbated if we observe progress in the creation of more and more versatile AIs. In contrast, they may reverse if one day software ceases to be the scarce input (and thus its share of output will go down), or if a well-designed global digital policy is implemented that addresses this issue.

6.2 Inequality Between Eras Trumps Inequality Within Eras

Let me now briefly comment on the impacts of the four technological revolutions on *inequality*, understood as variation within the distribution of certain variables across the human population. The key variables in question are income, wealth, or consumption of goods and services, though I expect that in the digital era it may also make sense to compute inequality in flows of useful data or access to computing power.

Many decompositions are useful for understanding the sources of inequality. Here I would like to concentrate on two of them: factor shares vs. concentration of factor ownership, and within-era vs. cross-era inequality.

The first of the above decompositions naturally corresponds to the aggregate production function approach I use in the current section. It can only be applied within a given era. The key logic is that changes in factor shares automatically bring changes in the degree of inequality as long as ownership of one factor is more concentrated than the other. And indeed ownership of "capital" K was typically much more concentrated than "labor" L, because agricultural labor, industrial labor, and human capital are naturally dispersed as they are embodied in humans. Land and physical capital, in contrast, tended to be concentrated in relatively few hands for centuries (Galor, 2005). The concentration of capital in the Western economies went down only in the early 20th century (Piketty and Zucman, 2014), during the world wars. In consequence the time path of inequality broadly followed the pattern of the capital share, but for the decline in inequality around both world wars.

The digital era seems to stand out from this analysis. For the first time in history, the "labor" factor L includes also disembodied pre-programmed software, whose ownership is highly concentrated. Therefore in the digital era increases in the "labor" L (software) share of output may possibly be even associated with increasing, not decreasing inequality, if bulk of the remuneration goes to computer software owners rather than human users. One should also be cautious when interpreting the landmark macroeconomic trends of the last 30 years: declining labor shares (Karabarbounis and Neiman, 2014), increasing profit shares (Barkai, 2017), and increasing inequality, especially at the top of the distribution (Piketty, 2014; Andrews, Criscuolo, and Gal, 2016; Jones and Kim, 2017). These observations are based on data which conflate the industrial-era and digital-era economy, and thus rather capture the shift between the eras and not illustrate any phenomena within the digital era. These findings are also confounded with the fact that technological revolutions hugely affect relative prices. After the industrial revolution there was a massive decline in the price of industrial goods relative to land and property. Similarly following the digital revolution we observe a decline in the price of ICT capital and other investment goods relative to consumption goods (and especially land and property), cf. Greenwood, Hercowitz, and Krusell (1997); Fernald (2015).

To understand the within-era vs. between-era inequality decomposition, one should begin with the observation that (i) technologies from subsequent eras coincide across the world and (ii) there are marked growth accelerations between eras. A large fraction of population in Sub-Saharan Africa is still employed in subsistence agriculture. At the same time, it becomes a continent of contrasts whenever we see a clash of the agricultural, industrial and digital eras. The Western world, in contrast, has largely eliminated subsistence agriculture, and achieved affluent and relatively equal societies. Yet, even there one may easily observe differences in the pace of development between the relatively stagnant traditional industry and the booming digital economy. This all suggests that the timing of takeoff into another era determines the relative standing of nations, regions, economic sectors, and people (Diamond, 1997). As growth rates can differ by orders of magnitude between eras, cross-era inequality quickly overwhelms within-era inequality, giving rise to a "great divergence" in economic development and bimodal (or more generally, multimodal) distributions of e.g., country or regional GDP (Quah, 1997; Henderson and Russell, 2005). Also within countries, measured increases in inequality, in particular top income inequality and CEO pay, may partly represent between-era inequality, i.e. the inequality between those who got on the train and those who missed it.

6.3 Measuring the Inputs in the Digital Era

Following my earlier call to design and implement Global Data Accounts in order to properly capture the digital-era output variable, *flows of useful data*, I would also like to call for better measurement of inputs of the digital-era production process, *hardware* and *software*. We also have to learn to measure hardware and software share of digital output, and digital inequality – by which I do not just mean inequality in available computing power, data storage capacity and bandwidth (akin to wealth inequality), but also in the in- and outflows of useful data (similar to income inequality).

Among the two inputs, measurement of hardware is probably going to be relatively easier. The "hardware" amalgamate may include all capital and consumption goods which have a digital component, but its key characteristics are likely going to be computing power, data storage capacity and bandwidth. First estimates have been already done (Hilbert and López, 2011). The "software" factor is going to be way more challenging, though. It may include the flows of services of existing (preprogrammed) software, but also the work of human programmers and other users. It could potentially be measured, e.g., as the aggregate number of elementary instructions provided to hardware. A caveat is that while human labor is rivalrous, software services are not because software can be costlessly copied and run simultaneously on many devices.

A related challenge would be to quantify hardware *capacity utilization*. Intuitively, hardware capacity utilization is likely generally low, much lower than physical capital utilization. We may likely find that computing power and bandwidth are idle most of the time, because our software (and particularly so, the human users) provide relatively very few instructions. This further strengthens my earlier assertion that huge returns are waiting for the developers of new software which could tap this unused potential and generate large flows of useful data.

There are also potentially interesting distributive implications from quantifying the volume of idle computing power and digital memory. Namely, their substantial fraction is probably going to be embodied in personal computers and other electronic devices whose ownership is dispersed. This raises a question if we could potentially implement a "computer banking" system allowing users to rent their idle computational capacity and digital memory, so that they could use their electronic devices as source of capital income. That would act to reduce inequality. On the software side, however, that would probably further reinforce development of sophisticated software replacing human work and thus increasing inequality. Thinking of striking a balance, we have to bear in mind, though, that the latter phenomenon will likely continue into the future anyway.

One final caveat is that there will be less and less scope for comparing the "wealth of nations" in the fully globalized digital era. The primary level of measurement should probably be gradually shifted from the country level to the company level, the global value chain (that is, value network) level, and the global level. Technological companies are already multi-national and almost footloose. For example, real GDP grew in Ireland in 2015 by 25.6%, against the European Union average of 2.3% (Eurostat data, in PPS). Why? Because a few large multi-national but hitherto US-based companies have either moved their productive assets to Ireland (e.g., Apple moved some of its valuable intellectual property assets such as copyrights and patents on the design and technology) or domiciled in Ireland by buying a smaller Irish-registered rival and "inverting" into an Irish corporate structure. They did so to save on taxes. It is however doubtful if this statistical fact had any bearing on the actual incomes and productivity of Irish citizens.

6.4 The Case for Better Digital Policy

Acknowledging that we have entered a new, digital era of economic development, offers one more (perhaps controversial) insight. Namely, it appears that compared to the agricultural and industrial economy, the digital economy is virtually unregulated. Our landmark economic policies and institutions have been designed with the industrial-era economy in mind. They are either inefficient when applied to a digital-era economy, or not applicable at all. At the same time, dedicated digital-era economic policies or institutions have not been devised yet.

Although such observations are rarely made in the public debate, they should not come as a surprise: we are arguably still at an early stage of the digital era, and after previous technological revolutions it also took time to develop suitable policies and institutions, which were developed bottom-up. When the agricultural revolution of the Neolithic period allowed for accumulation of food surpluses, this had anarchic outcomes until the state emerged and enforced hierarchical order, much more centralized than in the previous hunter-gatherer era (Diamond, 1997). Analogously, the industrial revolution unleashed fast-growing, laissez-faire industrial capitalism which was gradually regulated by installing worker rights, taxes, public education, health insurance, etc. Ultimately in the Western world this led to the welfare state with an increasingly high share of state budget in GDP. The industrial era arguably strengthened the centralized state in terms of tax collection and law enforcement compared to the agricultural-era powers of feudal rulers.

So far the digital-era economy is very loosely regulated compared to the industrial-era economy. Authorities are not, e.g., taxing data flows (which would be a digital-era equivalent of VAT), taxing software providers for their returns to collected useful data (which would be an equivalent of income tax), enforcing penalties for distributing harmful data, etc. Consequences are abound. For example, there exists a sizable "dark web", a digital-era representation of the shadow economy. Cyberpolicing is lagging behind cybercrime. Organized Internet-based misinformation actions are able to affect democratic voting outcomes and public sentiments but so far cannot be traced back to their original organizers. There is an economic arms race towards most versatile and adaptive AI technologies but there is very little research on AI safety. Insufficient regulation arguably also facilitates the formation of global monopolies, large-scale intrusion into privacy, and generally abusing the position of power on the side of large technological corporations. There is also a legitimate reason behind this last set of outcomes, however: unprecedented scalability of the digital economy. For example, software applications can be immediately deployed to the entire global market. In the digital era, start-up companies can grow very fast very quickly, and thus incumbent companies have strong incentives to collect as much data as possible in order to be able to fill the emerging market niches first.

Designing policies and institutions for the digital era requires to ask some important questions. First, are we able to trace and quantify the objects of interest, such as data returns to software and flows of useful data? Note that in the early years of the industrial revolution, when only agricultural-era regulations and policies were in place, physical capital and industrial output were also viewed as rather fluid, elusive concepts, not so easy to quantify and regulate as land ownership and crop yields; now we do this on a daily basis.

Second, the digital-era economy is global, and costs of "relocating" the business are historically small: it is much easier to re-route data traffic than to re-route goods or establish overseas businesses. Thus sensible digital policy should arguably be globally administered. But who should take care of this and would these institutions be publicly legitimized? Global institutions like the World Economic Forum clearly recognize the digital revolution (Schwab, 2016). In contrast, our mindsets tend to be adjusted to the local and national levels and our economics is still rooted in the industrial era. Perhaps time is needed until the world accepts that there could be legitimate authorities at the global level, even if specializing only in digital policy.

7 Side Effects of Development

If the driving force behind past development has been the human drive to maximize local control, we should not be surprised that it had many unintended side effects. For one thing, changes in natural ecosystems always have consequences reaching beyond the direct objectives which motivated these changes. Nature is a highly complex system, and humans cannot fully comprehend its complexity. We only build simplified models, so even with the best of our intentions there are always aspects of reality which are not accounted for.

We do not pursue the best of our intentions, though. We do not even try to maximize long-run social welfare, and we are not preoccupied with sustainability or intergenerational equity. What we do is maximize our local control, here and now. We do not account for the external effects beyond these narrow confines. In consequence, unintended side effects of our activities appear due to insufficient scale of analysis and suboptimal cooperation (bounded rationality), as well as insufficient treatment of long-run problems (myopia).[5]

We also do not have sympathy for other species. Successful expansion of control on behalf of humans has always been associated with declines in control on behalf of other species. The period of our flourishing is a period of mass extinctions: the Holocene extinction (Kolbert, 2014).

In the following subsection I will discuss the ecological side effects of the cumulative increases in human local control, observed over the four eras of development. I will argue that what is good for our local control, may have unintended negative consequences at the global scale. Next I will move to a discussion of the relationship between human control and development of AI. Finally I will comment on the biggest risk of all: the existential threat from artificial general intelligence (AGI).

7.1 Ecological Side Effects

Although one could be tempted to think that prehistorical hunter-gatherers lived in harmony with nature, the reality was one of a constant *war*. Our early ancestors

[5] As a partial defense of humankind, I should say that rationality of other species is even more bounded, and their time perspective is even shorter.

were able to conquer so many natural habitats across the world only because they eliminated the most important species which previously ruled them. The *homo sapiens* drove other hominins to extinction, and other species followed suit.

That happened *before* the humans learned to transform the natural habitats in their own favor. The following agricultural era witnessed much stronger ecological side effects of human activity. Transformation of diverse ecosystems into monocultures of wheat, corn, bananas, etc., dramatically reduced biodiversity. Farming was also a cause of massive deforestation and soil erosion. Moreover, as in the agricultural era people often relied on simplified, narrow diets based on few local food crops, they faced risks of malnutrition. Most dramatically, myopia and insufficient coordination sometimes caused episodes of famine or wars over scarce resources. Perhaps the clearest example here is the history of Easter Island collapse (de la Croix and Dottori, 2008).

Just as the positive effects of economic growth accelerated in the industrial era, so did the ecological side effects. The growth of cities and industrial zones replaced natural landscapes with seas of concrete. The growing manufacturing industry caused air, soil and water pollution, engulfing the cities in smog. It also created huge demand for exhaustible fossil resources: energy resources like oil, gas and coal, metal ores, and even rare chemical elements such as uranium. Some of them are expected to be depleted within the next decades. Increased greenhouse gas emissions induced global warming. Our completely transformed life environments brought civilizational diseases such as obesity or allergies. And development of nuclear weapons imposed on us a constant risk of a deadly nuclear war.

The digital revolution – similarly to the earlier technological revolutions – causes its own side effects without eliminating the previous ones. It brings civilizational diseases to an entirely new level, with increased incidences of Internet, social media or videogame addiction, or social phobia. It also bombards our brains with *digital pollution*, demanding them to navigate around Internet hate, fake news, and the rapid spread of socially harmful ideas such as anti-vaccination movements, anti-system agendas or actions reinforcing our undue prejudices and xenophobia. Furthermore, on the one hand it puts us under constant stress by raising the frequency of received e-mails, instant messages, and notifications. On the other hand it offers us instant gratification by flooding our brains with easy-to-digest but entirely unnecessary information, which reduces our span of attention and teaches our brain to avoid deeper thinking, a mechanism called "the shallows" (Carr, 2010). In sum, while the agricultural and industrial eras completely reshaped the relationships between the human and the environment, and affected our bodies, the digital era completely reshapes the relationships among humans, and affects our brains.

7.2 Artificial Intelligence and Human Control

While seeking to maximize our local control, humans built various devices which augmented their innate abilities. We carved sticks to reach further than our bare hands could reach, we wrote our ideas on paper so that we would not have to remember them later, we built cars so that we could move faster, and we built computers so that we would not have to invert that bloody four-by-four matrix by hand. As these devices got more and more sophisticated, we outsourced more and more of our actions. Sometimes these developments backfired on us, when we hurt ourselves with the stick, crashed the car into a tree, or realized that our job as a janitor has just been replaced with a surveillance camera. But generally these innovations increased the span of our control.

Artificial intelligence may put this logic in question. On the one hand, AI is just a set of optimization algorithms which have been created by human programmers, and which can only pursue goals set by human programmers. Seen in this way, it is just one more external device augmenting our abilities. On the other hand, it gets frightening when the algorithm provides answers that are far beyond our comprehension. For example, DeepMind AlphaZero exhibits vastly superhuman performance at chess an Go, and it learned that in less than a day, only by self-play, without using human knowledge.[6]

Extrapolating past trends, in particular observing the pace of growth in idle computing power, digital memory and the associated returns to software, I expect that in the future the "software" input will to an increasing degree include AI. Unfortunately, by getting more versatile and adaptive, it will also become more substitutable with human skilled labor. So far our software gradually automates routine tasks but is complementary to advanced human skills (Autor and Dorn, 2013; Frey and Osborne, 2013), part of a broader phenomenon of skill-biased technical change (Caselli and Coleman, 2006) and capital-skill complementarity (Krusell, Ohanian, Ríos-Rull, and Violante, 2000). However, when AI software will manage to automate also highly skilled jobs (potentially leaving people with only AI development and programming: tasks which are extremely skill-demanding and thus implementable only by a very narrow group of specialists), one may expect dire consequences for the labor share of output and income inequality. According to machine learning researchers (Grace, Salvatier, Dafoe, Zhang, and Evans, 2017): "AI will outperform humans in (...) translating languages (by 2024), writing high-school essays (by

[6]DeepMind's *AlphaGo Zero* AI algorithm has mastered the (extremely sophisticated) game of Go to a level allowing it to beat all human competitors as well as its previous incarnation *AlphaGo* (which already beat all human competitors) by 100-0. *AlphaGo Zero* was trained without supervision and without using human knowledge. Its performance suggests a *qualitative* improvement in the game play compared to the strategies invented by professional Go players over centuries (Silver, Schrittwieser, Simonyan, et al., 2017).

2026), driving a truck (by 2027), working in retail (by 2031), writing a bestselling book (by 2049), and working as a surgeon (by 2053). Researchers believe there is a 50% chance of AI outperforming humans in all tasks in 45 years and of automating all human jobs in 120 years."

Even worse than the distributional consequences are the effects for human control. AI algorithms are going to increasingly make decisions on our behalf and we are going to be less and less able to understand the premises for these decisions. The technologies developed so far are narrow AI algorithms which are confined to relatively limited action spaces. We can switch them off and do things our way. They do not yet have the capability to expand the space of actions which are searched in order to maximize the objective. Today's AlphaZero is vastly superhuman at chess and Go, but is not able to go from there to beat the stock market or master nanotechnology. It cannot push its constraints, self-rewrite, and go rule (let alone, navigate) the world. But we should not exclude the possibility that one day an artificial general intelligence (AGI) might be able to do that.

The quest for AGI opens the threat that machine intelligence will one day obtain the capability to search a *very broad* space of actions, while maximizing some objective which (even if to a tiny extent) opposes our human local control. By the logic of intelligence explosion via a cascade of recursive self-improvements, this AGI may develop itself to be orders of magnitude faster and more accurate than our human local control maximization process. We are, in the end, limited by the cognitive capacities of our brains and data communication abilities, using only speech and writing. Then, one day we may realize that the objectives we had programmed into the AGI have led to unintended consequences which cannot be corrected anymore: the AGI is now superintelligent and will always outsmart us, successfully protecting its original objective (Bostrom, 2012; Hanson and Yudkowsky, 2013). In this way, we may fall victim to our own dynamic inconsistency of actions: the AGI may get out of hand, giving the humankind a very narrow (if at all) time margin for reaction.

As an economist I expect that once AGI becomes technically possible, technological corporations will not hesitate to create it. Its development will certainly help dramatically increase the local control by *these few* humans who programmed the AGI, by making them rich, famous and powerful, at least until the unintended side effects show up. The economic incentives for developing AGI are clearly there: as I argued above, it is software, not hardware, which is now the bottleneck of the digital economy.

7.3 Existential Threat from AGI

To fully comprehend that development of AGI causes an *existential threat to humankind*, one has to first reiterate the underlying driving force behind development. Before the cognitive revolution, it was evolution, which aimed to maximize genetic fitness of species. Later, it was our human local control maximization process – a sub-routine of evolution, embedded in all organisms, which got out of hand in the case of humans because our frontal cortex development allowed us to acquire the theory of mind and exceptional social skills, and thus we have passed the threshold for collective knowledge accumulation. At this point, the human local control maximization process got out of hand owing to its sheer pace (faster by orders of magnitude than the pace of evolution), recursive self-improvement (thanks to knowledge accumulation), and greed for resources.

But our own maximization process also has plentiful sub-routines. Among them, we may want to build AGI, algorithms endowed with computer hardware, which not only do quick and error-free computations, but also learn without supervision like AlphaZero, navigate the world like Google self-driving cars, plug in for data and power when needed, self-rewrite their code when it is useful for the goal like the Gödel machine of Schmidhuber (2009b), and do this all at once. Among other processes, this process is particularly likely to get out of human control, for the same reasons as the human local control maximization process got out of evolution's control: much faster pace of digital computation, by orders of magnitude, recursive self-improvement of machine intelligence, and the AGI's greed for resources.[7]

Once we recognize this possibility, the key question becomes: what should be the *objective function* of the AGI that we intend to build? Can we *ascertain* that its pursuit will bring favorable outcomes for the humankind? I am purposefully referring only to the objective function here, because other characteristics of the optimization process, such as constraints or behavioral rules (such as Asimov's Three Laws of Robotics) can always be relaxed, reinterpreted or otherwise circumvented if the optimizer is powerful in pursuing its goal. I am also consciously emphasizing the requirement that we should be *certain* that the AGI will be friendly. Otherwise it will be an existential threat. The reason is that a superintelligent machine which follows a goal which is not *fully* aligned with our long-run collective well-being – importantly, not just our local control subject to bounded rationality and myopia – may exterminate humankind without giving us notice to stop the process (Good, 1965). This may happen even if the goal is seemingly innocuous, like in the Bostrom's (2003) thought experiment of a paperclip maximizer. Muehlhauser and Salamon (2012) suggest that human extinction will in fact be the *default* outcome of machine superintelligence unless "we first solve the problem of how to build an AI with a stable, desirable utility function". Warnings against this scenario have been voiced, among others, by influential intellectuals such as Stephen Hawking, Frank

[7] It is humbling to recall that the design of the human brain remains clearly suboptimal in a multiplicity of dimensions. As remarked by Stephen Hawking, there is no physical law precluding particles from being organized in ways that perform even more advanced computations than the arrangements of particles in human brains.

Wilczek, Bill Gates, and Elon Musk.

A key trap when thinking about the existential threat from AGI is antropomorphization (Bostrom, 2012). Our faulty intuition suggests that an AGI, before it kills us, should gain consciousness, self-awareness, ability to modify its goals, and that it should first intentionally turn rogue towards humankind. In fact none of these qualities is needed. AGI may just be a powerful optimizer with very few or no human features. More precisely, Bostrom (2012) argues for an *orthogonality thesis*: "Intelligence and final goals are orthogonal axes along which possible agents can freely vary. In other words, more or less any level of intelligence could in principle be combined with more or less any final goal." Furthermore, it is intelligence itself which causes the threat, and not the final goal, because regardless of the final goal, there is going to be likely emergent convergence of auxiliary goals of any AGI (Omohundro, 2008; Bostrom, 2012), following the *instrumental convergence thesis*. The drives towards self-preservation, goal-content integrity, efficiency (technological perfection), creativity (cognitive enhancement), and resource acquisition will follow from almost any conceivable final goal. Please note that all these auxiliary AGI goals are akin to the "local control" that humans would like to maximize; however the final goal on behalf of humans, whatever that is, is due to evolution whereas the AGI's final goal will be set by the human programmers.

In the end, intelligence – i.e., efficient cross-domain optimization (Yudkowsky, 2013) – has in fact always been an existential threat. Just as humans are an existential threat to all other, less intelligent species, AGI can become an existential threat to the humankind (and all other species, too) if it surpasses us in terms of intelligence. Or maybe even earlier, when it surpasses human aggregate cross-domain optimization power by extensively using abundant hardware. "The AI does not hate you, nor does it love you, but you are made out of atoms which it can use for something else." (Yudkowsky, 2008). The bottom line is that while deliberating on the possible economic and societal benefits of sophisticated, self-improving AI software, we should not forget that it can also be a double-edged sword. The predicted "singularity" may end up both as utopia and dystopia.

8 Conclusion

In sum, in this paper I proposed a synthetic, qualitative theory of economic growth and technological progress over the entire human history, i.e., across the four eras of development: the hunter-gatherer era, the agricultural era, the industrial era and the digital era. Across all the eras I considered the following themes: knowledge accumulation, economic growth, key factors of production and their mutual relation, inequality, and side effects of development.

The key novel elements of my theory are: (i) the proposition that the universal driving force of development can be summarized as a process of human local control maximization; (ii) the idea that each technological revolution opens up an entirely new dimension of economic growth which operates on top of previous ones, and thus measurement of development should be era-specific; (iii) identification of scale of operations in R&D and the knowledge depreciation rate as an important determinants of the pace of technological progress; (iv) the observation that each new era not only accelerates growth, but also positively feeds back on the previous era; (v) the observation that between-era inequality tends to trump within-era inequality.

Building on this synthetic theory and the analogies with previous eras, I drew new predictions of future developments in the digital era. I paid special attention to the observation that human skilled labor tends to be complementary to hardware but substitutable with software. I obtained new useful results here thanks to taking a broader perspective on the digital age than it has been typically done in the literature: namely, that of entire millennia of economic growth and technological progress. I also discussed the possibility that further development of sophisticated software will end up in building artificial general intelligence (AGI), and that the AGI will pose an existential threat to humankind.

There are many threads in this paper which ought to be continued. First and foremost, specialized measurement of digital-era inputs and output, in terms of data units and not dollar value, ideally leading to Global Data Accounts. Second, elaboration of the case for global digital policies and institutions. Third, quantification of the role of R&D capital in technological progress and economic growth in the industrial and digital eras. Fourth, detailed analysis of the multi-channeled impact of AI technologies on economic growth. Fifth, ongoing research on AI safety.

References

Acemoglu, D., and D. Autor (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4, chap. 12, pp. 1043–1171. Elsevier.

Acemoglu, D., and P. Restrepo (2016): "The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment," Working paper no. 22252, National Bureau of Economic Research.

Andrews, D., C. Criscuolo, and P. N. Gal (2016): "The Global Productivity Slowdown, Technology Divergence and Public Policy: A Firm Level Perspective," Working party no. 1 on macroeconomic and structural policy analysis, OECD.

Ashraf, Q., and O. Galor (2013): "The 'Out of Africa' Hypothesis, Human Genetic Diversity, and Comparative Economic Development," *American Economic Review*, 103, 1–46.

Autor, D. H., and D. Dorn (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103(5), 1553–97.

Barkai, S. (2017): “Declining Labor and Capital Shares,” Job market paper, University of Chicago.

Barro, R. J., and X. X. Sala-i-Martin (2003): *Economic Growth*. MIT Press.

Bertocchi, G. (2006): “The Law of Primogeniture and the Transition from Landed Aristocracy to Industrial Democracy,” *Journal of Economic Growth*, 11, 43–70.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2017): “Are Ideas Getting Harder to Find?,” Working paper, Stanford University.

Bocquet-Appel, J.-P. (2011): “When the World’s Population Took Off: The Springboard of the Neolithic Demographic Transition,” *Science*, 333, 560–561.

Bostrom, N. (2003): “Ethical Issues in Advanced Artificial Intelligence,” in *Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence, Vol. 2*, ed. by G. E. Lasker, W. Wallach, and I. Smit, pp. 12–17. International Institute of Advanced Studies in Systems Research and Cybernetics.

——— (2012): “The Superintelligent Will: Motivation and Instrumental Rationality in Advanced Artificial Agents,” *Minds and Machines*, 22, 71–85.

Boucekkine, R., D. de la Croix, and O. Licandro (2003): “Early Mortality Declines at the Dawn of Modern Growth,” *Scandinavian Journal of Economics*, 105, 401–418.

Bowlby, J. (1969): *Attachment*. The Hogarth Press and The Institute of Psychoanalysis.

Brynjolfsson, E., and A. McAfee (2014): *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W.W. Norton & Co.

Carr, N. G. (2010): *The Shallows: What the Internet Is Doing to Our Brains*. W. W. Norton & Company.

Caselli, F., and W. J. Coleman (2006): “The World Technology Frontier,” *American Economic Review*, 96, 499–522.

Chirinko, R. S. (2008): “σ: The Long and Short of It,” *Journal of Macroeconomics*, 30, 671–686.

Comin, D., and B. Hobijn (2010): “An Exploration of Technology Diffusion,” *American Economic Review*, 100, 2031–2059.

de la Croix, D., and D. Dottori (2008): "Easter Island's Collapse: A Tale of a Population Race," *Journal of Economic Growth*, 13, 27–55.

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

Dunbar, R. I. M. (1992): "Neocortex Size as a Constraint on Group Size in Primates," *Journal of Human Evolution*, 22, 469–493.

——— (1993): "Coevolution of Neocortex Size, Group Size, and Language in Humans," *Behavioral and Brain Sciences*, 16, 681–735.

Dunsworth, H. M. (2010): "Origin of the Genus Homo," *Evolution: Education and Outreach*, 3, 247.

Felipe, J., and F. M. Fisher (2003): "Aggregation in Production Functions: What Applied Economists Should Know," *Metroeconomica*, 54, 208–262.

Fernald, J. G. (2015): "Productivity and Potential Output Before, During, and After the Great Recession," *NBER Macroeconomics Annual*, 29, 1–51.

Frey, C. B., and M. Osborne (2013): *The Future of Employment: How Susceptible Are Jobs to Computerisation?* Oxford Martin Programme on the Impacts of Future Technology.

Gabaix, X., and A. Landier (2008): "Why Has CEO Pay Increased So Much?," *Quarterly Journal of Economics*, 123, 49–100.

Galor, O. (2005): "From Stagnation to Growth: Unified Growth Theory," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. N. Durlauf, pp. 171–293. North-Holland.

——— (2011): *Unified Growth Theory.* Princeton: Princeton University Press.

Gillings, M. R., M. Hilbert, and D. J. Kemp (2016): "Information in the Biosphere: Biological and Digital Worlds," *Trends in Ecology and Evolution*, 31, 180–189.

Good, I. J. (1965): "Speculations Concerning the First Ultraintelligent Machine," *Advances in Computers*, 6, 31–88.

Gordon, R. J. (2016): *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War.* Princeton University Press.

Grace, K., J. Salvatier, A. Dafoe, B. Zhang, and O. Evans (2017): "When Will AI Exceed Human Performance? Evidence from AI Experts," Working paper, Future of Humanity Institute, Oxford University.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997): "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 87, 342–362.

Growiec, J. (2007): "Beyond the Linearity Critique: The Knife-Edge Assumption of Steady-State Growth," *Economic Theory*, 31, 489–499.

Growiec, J., and I. Schumacher (2013): "Technological Opportunity, Long-Run Growth, and Convergence," *Oxford Economic Papers*, 65, 323–351.

Growiec, K., and J. Growiec (2014): "Trusting Only Whom You Know, Knowing Only Whom You Trust: The Joint Impact of Social Capital and Trust on Happiness in CEE Countries," *Journal of Happiness Studies*, 15, 1015–1040.

Ha, J., and P. Howitt (2007): "Accounting for Trends in Productivity and R&D: A Schumpeterian Critique of Semi-Endogenous Growth Theory," *Journal of Money, Credit and Banking*, 39(4), 733–774.

Hansen, G. D., and E. C. Prescott (2002): "Malthus To Solow," *American Economic Review*, 92, 1205–1217.

Hanson, R. (2000): "Long-Term Growth as a Sequence of Exponential Modes," Working paper, George Mason University.

Hanson, R., and E. Yudkowsky (2013): *The Hanson-Yudkowsky AI-Foom Debate*. Machine Intelligence Research Institute.

Harari, Y. N. (2014): *Sapiens: A Brief History of Humankind*. Vintage.

——— (2017): *Homo Deus: A Brief History of Tomorrow*. Vintage.

Henderson, D. J., and R. R. Russell (2005): "Human Capital and Convergence: A Production–Frontier Approach," *International Economic Review*, 46, 1167–1205.

Hibbs, D. A., and O. Olsson (2004): "Geography, Biogeography, and Why Some Countries Are Rich and Others Are Poor," *Proceedings of the National Academy of Sciences*, 101, 3715–3720.

Hilbert, M. (2017): "The More You Know, the More You Can Grow: An Information Theoretic Approach to Growth in the Information Age," *Entropy*, 19, 82.

Hilbert, M., and P. López (2011): "The World's Technological Capacity to Store, Communicate, and Compute Information," *Science*, 332, 60–65.

Johansen, A., and D. Sornette (2001): "Finite-time Singularity in the Dynamics of the World Population, Economic and Financial Indices," *Physica A*, 294, 465–502.

Jones, B. F. (2009): "The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder? ," *Review of Economic Studies*, 76, 283–317.

Jones, C. I. (1995): "R&D-Based Models of Economic Growth," *Journal of Political Economy*, 103, 759–84.

——— (1999): "Growth: With or Without Scale Effects?," *American Economic Review*, 89(2), 139–144.

——— (2005): "Growth and Ideas," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. Durlauf. North-Holland.

Jones, C. I., and J. Kim (2017): "A Schumpeterian Model of Top Income Inequality," *Journal of Political Economy*, p. forthcoming.

Jones, C. I., and P. M. Romer (2010): "The New Kaldor Facts: Ideas, Institutions, Population, and Human Capital," *American Economic Journal: Macroeconomics*, 2, 224–245.

Jorgenson, D. W. (2005): "Accounting for Growth in the Information Age," in *Handbook of Economic Growth*, ed. by P. Aghion, and S. Durlauf. North-Holland.

Jorgenson, D. W., and K. J. Stiroh (2000): "Raising the Speed Limit: U.S. Economic Growth in the Information Age," *Brookings Papers on Economic Activity*, 1, 125–211.

Kaldor, N. (1961): "Capital Accumulation and Economic Growth," in *The Theory of Capital*, ed. by L. A. Lutz, and D. C. Hague, pp. 177–222. Palgrave Macmillan.

Karabarbounis, L., and B. Neiman (2014): "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 129(1), 61–103.

Klump, R., P. McAdam, and A. Willman (2007): "Factor Substitution and Factor Augmenting Technical Progress in the US," *Review of Economics and Statistics*, 89, 183–192.

——— (2012): "Normalization in CES Production Functions: Theory and Empirics," *Journal of Economic Surveys*, 26, 769–799.

Kolbert, E. (2014): *The Sixth Extinction: An Unnatural History*. Henry Holt and Company.

Kremer, M. (1993): "Population Growth and Technological Change: One Million B.C. to 1990," *Quarterly Journal of Economics*, 108, 681–716.

Kruse-Andersen, P. K. (2017): "Testing R&D-Based Endogenous Growth Models," Working paper, University of Copenhagen.

Krusell, P., L. E. Ohanian, J.-V. Ríos-Rull, and G. L. Violante (2000): "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68, 1029–1054.

Kurzweil, R. (2005): *The Singularity is Near*. New York: Penguin.

Liu, H., F. Prugnolle, A. Manica, and F. Balloux (2006): "A Geographically Explicit Genetic Model of Worldwide Human-Settlement History," *American Journal of Human Genetics*, 79, 230–237.

Liu, W., M. Martinón-Torres, Y.-j. Cai, S. Xing, H.-w. Tong, S.-w. Pei, M. J. Sier, X.-h. Wu, R. L. Edwards, H. Cheng, Y.-y. Li, X.-x. Yang, J. M. Bermúdez de Castro, and X.-j. Wu (2015): "The Earliest Unequivocally Modern Humans in Southern China," *Nature*, 526, 696–699.

Madsen, J. (2008): "Semi-endogenous Versus Schumpeterian Growth Models: Testing the Knowledge Production Function Using International Data," *Journal of Economic Growth*, 13, 1–26.

Maslow, A. (1954): *Motivation and Personality*. Harper & Brothers.

Mokyr, J. (2002): *The Gifts of Athena: Historical Origins of the Knowledge Economy*. Princeton University Press.

Muehlhauser, L., and A. Salamon (2012): "Intelligence Explosion: Evidence and Import," in *Singularity Hypotheses: A Scientific and Philosophical Assessment*, ed. by A. Eden, J. Søraker, J. H. Moor, and E. Steinhart. Berlin: Springer.

Nagy, B., J. D. Farmer, J. E. Trancik, and J. P. Gonzales (2011): "Superexponential Long-Term Trends in Information Technology," *Technological Forecasting and Social Change*, 78, 1356–1364.

Oberfield, E., and D. Raval (2015): "Micro Data and Macro Technology," *Econometrica*, p. forthcoming.

Olsson, O. (2000): "Knowledge as a Set in Idea Space: An Epistemological View on Growth," *Journal of Economic Growth*, 5, 253–275.

——— (2005): "Technological Opportunity and Growth," *Journal of Economic Growth*, 10, 31–53.

Omohundro, S. M. (2008): "The Basic AI Drives," in *Proceedings of the First AGI Conference, Volume 171, Frontiers in Artificial Intelligence and Applications*, ed. by P. Wang, B. Goertzel, and S. Franklin. IOS Press.

Piketty, T. (2014): *Capital in the Twenty-First Century*. Harvard University Press, Translated from French by Arthur Goldhammer.

Piketty, T., and E. Saez (2003): "Income Inequality in the United States, 1913–1998," *Quarterly Journal of Economics*, 118.

Piketty, T., and G. Zucman (2014): "Capital is Back: Wealth-Income Ratios in Rich Countries, 1700-2010," *Quarterly Journal of Economics*, 129(3), 1255–1310.

Quah, D. T. (1997): "Empirics for Growth and Distribution: Stratification, Polarization, and Convergence Clubs," *Journal of Economic Growth*, 2, 27–59.

Romer, P. M. (1986): "Increasing Returns and Long-run Growth," *Journal of Political Economy*, 94(5), 1002–1037.

——— (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98, S71–S102.

Schmidhuber, J. (2009a): "Driven by Compression Progress: A Simple Principle Explains Essential Aspects of Subjective Beauty, Novelty, Surprise, Interestingness, Attention, Curiosity, Creativity, Art, Science, Music, Jokes," in *Anticipatory Behavior in Adaptive Learning Systems, from Sensorimotor to Higher-level Cognitive Capabilities*, ed. by G. Pezzulo, M. V. Butz, O. Sigaud, and G. Baldassarre. Springer, LNAI.

——— (2009b): "Ultimate Cognition à la Gödel," *Cognitive Computation*, 1, 177–193.

Schwab, K. (2016): *The Fourth Industrial Revolution*. World Economic Forum.

Silver, D., J. Schrittwieser, K. Simonyan, et al. (2017): "Mastering the Game of Go Without Human Knowledge," *Nature*, 550, 354–359.

Solow, R. M. (1957): "Technical Change and the Aggregate Production Function," *Review of Economics and Statistics*, 39, 312–320.

——— (1987): "We'd Better Watch Out," *New York Times Book Review*, July 12, 36.

Tattersall, I. (2009): "Human Origins: Out of Africa," *Proceedings of the National Academy of Sciences*, 106, 16018–16021.

TEMPLE, J. (2006): "Aggregate Production Functions and Growth Economics," *International Review of Applied Economics*, 20, 301–317.

TIMMER, M. P., AND B. VAN ARK (2005): "IT in the European Union: A driver of productivity divergence?," *Oxford Economic Papers*, 57(4), 693–716.

TOMASELLO, M. (2014): *A Natural History of Human Thinking*. Harvard University Press.

TWENGE, J. M., AND W. K. CAMPBELL (2009): *The Narcissism Epidemic: Living in the Age of Entitlement*. Simon & Schuster, Inc.

WISSNER-GROSS, A. D., AND C. E. FREER (2013): "Causal Entropic Forces," *Physical Review Letters*, 110, 168702.

YUDKOWSKY, E. (2008): "Artificial Intelligence as a Positive and Negative Factor in Global Risk," in *Global Catastrophic Risks*, ed. by N. Bostrom, and M. M. Ćirković, p. 308–345. New York: Oxford University Press.

YUDKOWSKY, E. (2013): "Intelligence Explosion Microeconomics," Technical report 2013-1, Machine Intelligence Research Institute.