Ernst, Ekkehardt; Merola, Rossana; Samaan, Daniel

Article

Economics of artificial intelligence: Implications for the future of work

IZA Journal of Labor Policy

Provided in Cooperation with:
IZA – Institute of Labor Economics

Suggested Citation: Ernst, Ekkehardt; Merola, Rossana; Samaan, Daniel (2019) : Economics of artificial intelligence: Implications for the future of work, IZA Journal of Labor Policy, ISSN 2193-9004, Sciendo, Warsaw, Vol. 9, Iss. 1, pp. 1-35, https://doi.org/10.2478/izajolp-2019-0004

This Version is available at:
https://hdl.handle.net/10419/222165

Ekkehardt Ernst, Rossana Merola and Daniel Samaan*

Economics of Artificial Intelligence: Implications for the Future of Work

Abstract

The current wave of technological change based on advancements in artificial intelligence (AI) has created widespread fear of job loss and further rises in inequality. This paper discusses the rationale for these fears, highlighting the specific nature of AI and comparing previous waves of automation and robotization with the current advancements made possible by a widespread adoption of AI. It argues that large opportunities in terms of increases in productivity can ensue, including for developing countries, given the vastly reduced costs of capital that some applications have demonstrated and the potential for productivity increases, especially among the low skilled. At the same time, risks in the form of further increases in inequality need to be addressed if the benefits from AI-based technological progress are to be broadly shared. For this, skills policies are necessary but not sufficient. In addition, new forms of regulating the digital economy are called for that prevent further rises in market concentration, ensure proper data protection and privacy, and help share the benefits of productivity growth through the combination of profit sharing, (digital) capital taxation, and a reduction in working time. The paper calls for a moderately optimistic outlook on the opportunities and risks from AI, provided that policymakers and social partners take the particular characteristics of these new technologies into account.

Current version: February 07, 2019
Keywords: artificial intelligence, technological unemployment, inequality, productivity growth
JEL-codes: J23, J24, O00, E24
Corresponding author: Ekkehardt Ernst
ernste@ilo.org

Research Department, International Labour Organisation, Route des Morillons 4, 1211 Geneva, Switzerland. E-mail: ernste@ilo.org; merola@ilo.org; samaan@ilo.org

Cite as: Ernst et al. *IZA Journal of Labor Policy* (2019) 9:4.
https://doi.org/10.2478/izajolp-2019-0004

> My research activities during the past decade have brought me in contact with developments in the use of electronic digital computers. These computers are startling even in a world that takes atomic energy and prospects of space travel in its stride. The computer and the new decision-making techniques associated with it are bringing changes in white-collar, executive, and professional work as momentous as those the introduction of machinery has brought to manual jobs. (Simon, 1960)

1 Introduction

Values, norms, and language have evolved over the last six decades. What has remained the same, however, is the fear of the machine. Herbert Simon, Nobel Prize winner in economics, expressed in 1956 what many observers were convinced of at the time: "Machines will be capable, within twenty years, of doing any work a man can do," and hence that new technologies would make many jobs obsolete beyond the traditional blue-collar work in the manufacturing sweatshops. Today, we have grown used to computers around us: at home, in the office, at the bank, when travelling, or simply ordering food at the next drive-in restaurant. Rarely do we think of the jobs that might have been lost because of these computers and machines. Today, we no longer fear the computer that Professor Simon was afraid of, but something more profound: artificial intelligence (AI) or the capacity of machines to make predictions using large amounts of data to take actions in complex, unstructured environments (Agrawal et al., 2018a).

Complex decision-making under uncertainty is at the heart of modern economies. Whether as a consumer deciding which products and services to consume, as an employee when it comes to choosing the right job and career, or as a manager when running daily operations or planning the next factory, we all face constantly and simultaneously complex, interrelated problems for which our natural intelligence seems to have made us particularly well equipped. Indeed, until recently, no machines were remotely deemed to be capable of matching our intellectual capacity, even though the idea of an intelligent machine emerged as soon as the invention of the computer in the 1930s. In 1936, long before the invention of modern, silicon-based computers, Alonzo Church and Alan Turing – independently from each other – discovered that any process of formal reasoning – such as problems in economics and management described above – can be simulated by digital machines. In other words, the difference between a computer and a brain is one in degree, not in principle. Turing (1950) later argued that there might be a time when humans would no longer be able to distinguish between interacting with another human or a digital machine, passing the so-called "Turing test". Moreover, indeed, in light of recent experiences by leading AI firms, this time no longer seems to be too far away.

Intelligent digital assistants such as the "Google Assistant" which can be assigned to autonomously make appointments over the phone is but one possible application of AI (OECD, 2017).[1] Speech and image recognition, natural language processing, and machine translation figure prominently as key areas of development around AI. Others include automatic text generation such as the preparation of (short) journalistic pieces, automatic generation of company statements, or customer tele-assistants. More sophisticated applications include medical expert systems to analyze and diagnose patients' pathologies (medtech), automated review of legal

1 See https://www.youtube.com/watch?v=ogfYd705cRs&t=2150s for a short demonstration.

contracts to prepare litigation cases (lawtech),[2] self-driving cars or trucks, and the detection of patterns in stock markets for successful trading (algorithmic trading). Even creative arts, an area supposedly specific to humans, has seen a proliferation of applications in AI, from computers composing new pieces of music to painting programs replicating pictures in the style of a Rembrandt.[3] Common to all these applications is that they concern tasks that are considered to require specific human capacities related to visual perception, speech, sentiment recognition, and decision-making. In other words, AI is replacing mental tasks rather than physical ones, which were the target of previous waves of mechanization.

These advancements in AI have been made possible thanks to the confluence of three different, albeit related developments:

- A phenomenal drop in computing costs has led to an explosion in installed computing power and storage capacity. Simple smartphones today are significantly more powerful than the computer that brought the first man to the moon. The costs for producing an iPhone 7, for instance, currently stands at around US$220; in the 1980s, it would have been around US$1.2 million in today's terms simply to pay for the memory capacity of such a phone.
- Second, the development and widespread adoption of the Internet and other forms of digital communication has led to a significant increase in the supply and storage of digital information, including in central locations (cloud computing), which allow the comparison and analysis of significant amounts of data for statistical purposes that are necessary to develop tools based on AI principles.
- Finally, the drop in capital costs for digital technologies has significantly lowered barriers of entry for start-ups, making it less necessary than in the past to mobilize huge amounts of capital before starting a new venture while at the same time offering substantial first-mover advantages. This shift in business models toward small, rapidly growing tech companies was often driven by university spin-offs funded through innovative financial products and supported by a seemingly endless supply of highly educated software engineers. A paradoxical consequence of the digital nature of latest innovations is that the lower barriers to entry have allowed new players to uproot incumbents while at the same time quickly leading to new forms of industry concentration (Bessen, 2017a).

Together, these three developments triggered a rapid increase in AI patent applications across different patent offices worldwide (Fig. 1). As a result, an endless stream of new services and products appeared, with those surviving the test of the market growing rapidly in size and quickly overtaking large, well-established companies in traditional business lines. Indeed, within the short period of 15 years, companies such as Google, Apple, Facebook, and Amazon have belittled historic behemoths of American capitalism of the likes of Walmart, General Motors, or General Electric.

This sudden burst in applications of AI has created the sentiment of vastly accelerating technological change that is feared to disrupt labor markets in yet unforeseen magnitude (Ernst, 2018). What is puzzling, however, is that so far and despite the apparent acceleration in

2 For more examples from the legal industry, see https://law-tech-a2j.org/ai/artificial-intelligence-legal-services-and-justice/.

3 See https://www.livescience.com/54364-computer-creates-new-rembrandt-painting.html.

Figure 1 Number of AI patents granted by country

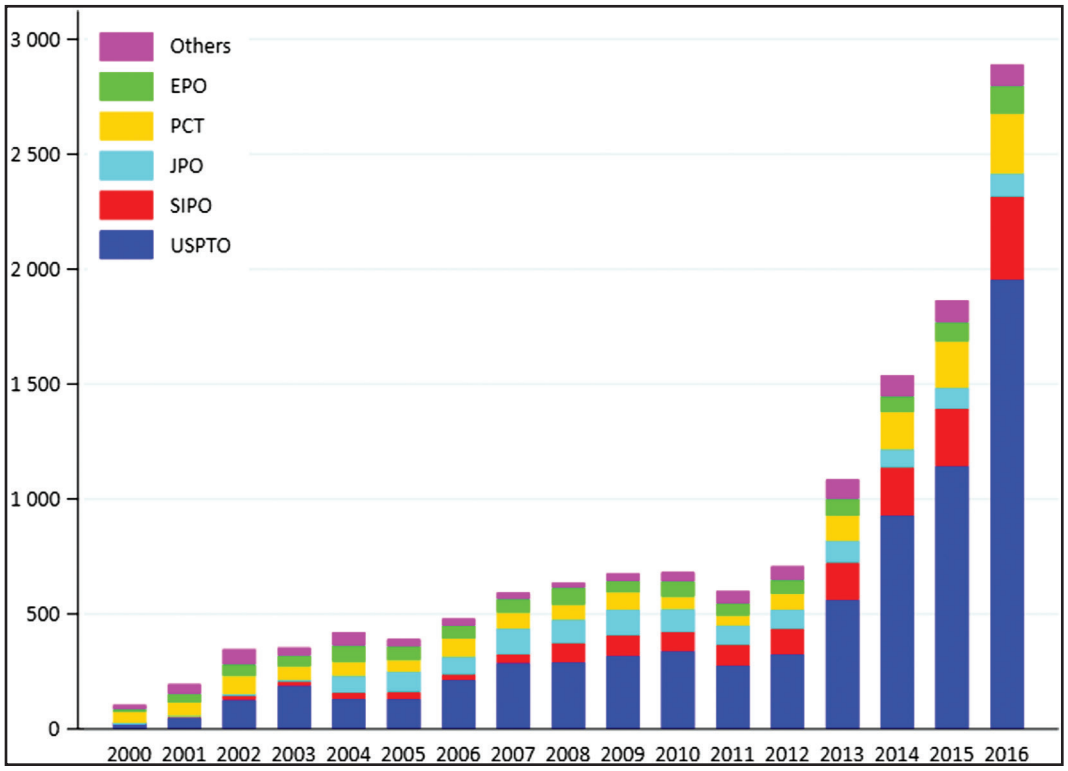

Source: Fujii and Managi, 2018. AI, artificial intelligence; EPO, European Patent Office; JPO, Japan Patent Office; PCT, Patent Cooperation Treaty; SIPO, State Intellectual Property Office of the People's Republic of China; USPTO, United States Patent and Trademark Office.

technological change, productivity growth has continued to decline in advanced economies. Similarly, no disruption seems to have struck global labor markets so far which, on the contrary, seem to have recovered from their slump following the global financial crisis (ILO, 2018). What has changed is a continuous worsening of country-level income inequality, continuing a long-term trend that started in the 1980s. But even here, looking at the global level where inequality and poverty rates have fallen thanks largely to emerging economies catching up, neither the expected benefits nor feared costs of automation – even less of AI – have yet materialized at a large scale.

Most observers are not reassured, however. Many analysts are warning that advances in both robotics and AI over the next few decades could lead to significant job losses or job polarization and hence widen income and wealth disparities (Korinek and Stiglitz, 2017; Méda, 2016). A recent report by Bank of America Merrill Lynch in 2015 pointed to the potential for a rise in inequality as a result of increased automation. The report cited research by Oxford University, which found that up to 35% of all workers in the United Kingdom, and 47% of those in the United States, are at risk of being displaced by technology over the next 20 years (Frey and Osborne, 2017). According to the World Bank (2016), in developing countries many more jobs are at risk: 69% in India, 72% in Thailand, 77% in China, and a massive 85% in Ethiopia. Other researchers, however, reach much less dramatic conclusions (Arntz et al., 2016, 2017). Nevertheless, what all these studies have in common is that they focus on *potential* gross job destruction and cannot provide an answer to *actual* job destruction, net job displacements, or labor market turnover, which would be necessary to assess the challenge

of automation from a policy perspective. Moreover, it is unclear to what extent conclusions can be drawn from many of the existing studies for technologies such as AI, on which little is known and almost no data exist.

This paper aims at addressing this knowledge gap to gain a better understanding of the economic and social implications of AI. To do so, it suggests starting from a granular analysis of how previous waves of automation have changed occupations and employment opportunities in the past. Specifically, we look at experiences of advanced and emerging economies with the automation of *physical* tasks through the rise in robotization. This approach can shed some light on the likely impact that the development and widespread diffusion of AI might have on employment, incomes, and inequality through the automation of *mental* tasks – as per our distinction between AI and robots/mechanization above. We also look at offshoring, in as much as it affects the role that AI can play in the structural transformation in developing countries. The paper then tries to answer the following questions. First, to what extent is the current digital transformation through the rise in AI labor augmenting rather than labor saving? Moreover, what will be the implications for productivity and inequality given the specific, digital nature of AI applications? In particular, can we expect an acceleration in productivity and earnings growth thanks to widespread diffusion of AI in areas that have not yet been subject to large-scale automation? Or, on the contrary, should we be afraid of technological rents arising from AI to be appropriated by the lucky few?

The answer that this paper gives to these questions is moderately optimistic. New, AI-based digital technologies may allow larger segments of the labor market to improve their productivity and to access better paying occupations and, thereby, may help promote (inclusive) growth. This requires, however, that a certain number of policies are put in place that support the necessary shift in occupational demand, maintain a strong competitive environment to guarantee diffusion of innovation, and keep up aggregate demand to support structural transformation. At the same time, AI applications raise the potential for productivity growth for interpersonal, less technical occupations and tasks, leading to higher demand for such work, which is likely to dampen the inequality trends observed over recent decades. A particular challenge arises for developing countries when they are part of a supply chain that forces them to adopt capital-intensive technologies despite an abundance of underutilized labor. Here, AI-driven automation might further drive up informality unless governments ensure a widespread adoption and diffusion of digital technological change beyond the supply chain sectors. In other words, the productivity-enhancing potential of AI is real but the specific characteristics of this new technology require policy responses that differ from those given during previous waves of technological change to generate shared benefits for the world of work.

To develop our argument, this paper starts with a historical perspective on automation. It argues that the rise in educational attainment has led to an increasing skill-biased nature of technological change, bringing fewer benefits for productivity but increasing inequality; it is against this background that the introduction of AI needs to be assessed. Section 3 shifts the focus on tasks and away from jobs to help understand the implications this has had for employment and the organization of production, before Section 4 discusses the particular experience that advanced and emerging economies have made during the recent wave of robotization. In Section 5, our focus then turns toward AI and the various effects on job growth, earnings dynamics, and firm productivity. In Section 6, we develop possible policy answers that can

help address the issues that AI brings to allow for a proper sharing of technological rents both within countries and between advanced and less developed economies.

2 Automation and productivity in historical perspective

Historically, productivity and living standards have increased thanks to a continuous division of labor (specialization) and replacement of more tedious, arduous, and routine tasks by machines. In agriculture, for instance, a modern farmer buys sophisticated machinery for the industrial production of farm goods to be sold through regional distribution centers, rather than using self-made tools to plow one's acre for self-consumption as it was done for centuries. Highly specialized labor at each level of such supply chains that work through automated processes allows for a timely production of goods and services at constant, predefined levels of quality and quantity. Moreover, agriculture was only the first sector to benefit from automation, given its dominance even in advanced economies until the 1950s in terms of total number of jobs. Thanks to the invention of the steam mill and later to widespread electrification, manufacturing of goods from textile to automobiles pushed out the boundaries of productivity thanks to a combination of automation and ever finer division of labor.

In contrast to fears expressed today, the wave of automation that came with the first and second industrial revolution during the 19th and early 20th centuries led to a rapid increase in demand for low- or unskilled labor, raising concerns about the demeaning nature of technological change (Braverman, 1974; Marglin, 1974).[4] As productivity growth in agriculture led to a massive shedding of labor in this sector, unskilled laborers often found new employment opportunities in manufacturing or other sectors such as mining and construction that were blossoming thanks to automation. As the division of labor progressed, workers were asked to concentrate on ever narrower, highly repetitive tasks to be performed at high speed. This so-called Taylorist approach to the organization of work – also dubbed a "scientific management" approach to production by organizational specialist Frederick Taylor – created significant strain among workers who were less and less able to identify themselves with the final outcome of their work. As a consequence, in the 1960s, social movements started to flare up to express demand for less demeaning work, better working conditions, and faster wage growth. At the same time, this was also the moment when productivity growth was highest among advanced economies, lifting large parts of the population out of poverty and creating a quickly expanding middle class.

With the rise in income, educational attainment grew as well. As (young) workers became increasingly educated, technological change shifted gears, laying the ground for the advent of the third industrial revolution based on the introduction of computers (Acemoglu, 2002). In the decades following the 1970s, technological change became skill biased, increasing gradually the demand for medium- and high-skilled workers at the expense of those with only primary education levels or less. Although the observed rise in unemployment was only partly technological and in large part driven by changes in the macroeconomic environment, work

4 There is no commonly accepted classification of different stages of industrial advancements. The notion currently most in use is to talk about artificial intelligence and related innovations as the Fourth Industrial Revolution, see Schwab (2016). Previous stages include the introduction of the steam engine first industrial revolution (IR), the widespread use of electricity (second IR), and the use of computers (third IR).

processes started to change, with manufacturing employment falling gradually in all major advanced economies as more and more sophisticated machines and robots – "automatically controlled, reprogrammable, multipurpose manipulator, programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications" (International Organization for Standardization, ISO) – would replace routine and repetitive tasks. At the same time, designing, implementing, and maintaining these robots and computers led to the emergence of a whole new industry, albeit offering significantly less employment opportunities than those lost in the process of automation. Overall, existing studies suggest that employment effects specifically from the introduction of robots remained rather limited or – depending on the methodology used – were even positive in the aggregate (Acemoglu and Restrepo, 2017; Bessen, 2017b; Chiaccio et al., 2018; De Backer et al., 2018, Graetz and Michaels, 2015). When extending the analysis to developing countries, however, the introduction of robots shows significant and much more substantial negative effects on employment (Carbonero et al., 2018).

With the decline in manufacturing employment, the service sector took over the role of a jobs engine. Business services, transportation, and distribution (wholesale and retail) among others offered new jobs tailored to better educated and trained people in the workforce. From the 1990s onward, concerns over automation were limited to a smaller and smaller workforce in manufacturing and attention shifted toward working conditions and opportunities in services. In particular in the United States, the advent of information and communication technologies (ICTs) produced a boom in investment in new technologies, which accelerated – temporarily – productivity growth and offered new employment opportunities, albeit often under less favorable conditions than what had been experienced during the boom in manufacturing employment.

Nevertheless, this third industrial revolution based on ICT innovations and the introduction of robots has brought much less economic benefits than the previous two waves of technological change. Indeed, looking at economic development in seven selected leading economies from a long-term perspective, a deceleration in productivity growth can be detected, despite a short-lived acceleration during the 1990s (Fig. 2). This is also reflected in similar developments of Gross Domestic Product (GDP) per capita (i.e., including the inactive population) that show a remarkable absence of accelerating improvements in living standards. This observation had already perplexed economists during the 1980s when Robert Solow famously stated that "you can see the computer age everywhere but in the productivity statistics" (David, 1990). Besides measurement issues related to the digital nature of ICT innovations, this might be related to the fact that improvements in ICT impacted only a few sectors (notably transportation and logistics industries besides telecommunications) in contrast to previous, general purpose technologies such as electricity (Gordon, 2016).[5]

The introduction of robots also offered new opportunities for automation along global supply chains, triggering a flourishing discussion about the global employment effects of off- and re-shoring in both developed and emerging economies. The United Nations Conference on Trade and Development (UNCTAD, 2016) argues that the historical labor cost advantage of low-income countries might be eroded by robots if they become cheap and easily substitutable

5 The debate on the slowdown in measured productivity growth has been an active field of research in recent years and goes beyond the scope of this paper. For an overview of the different arguments, see https://www.brookings.edu/research/the-productivity-slump-fact-or-fiction-the-measurement-debate/.

Figure 2 Hourly productivity growth rates in Group of Seven (G7) countries, decade averages

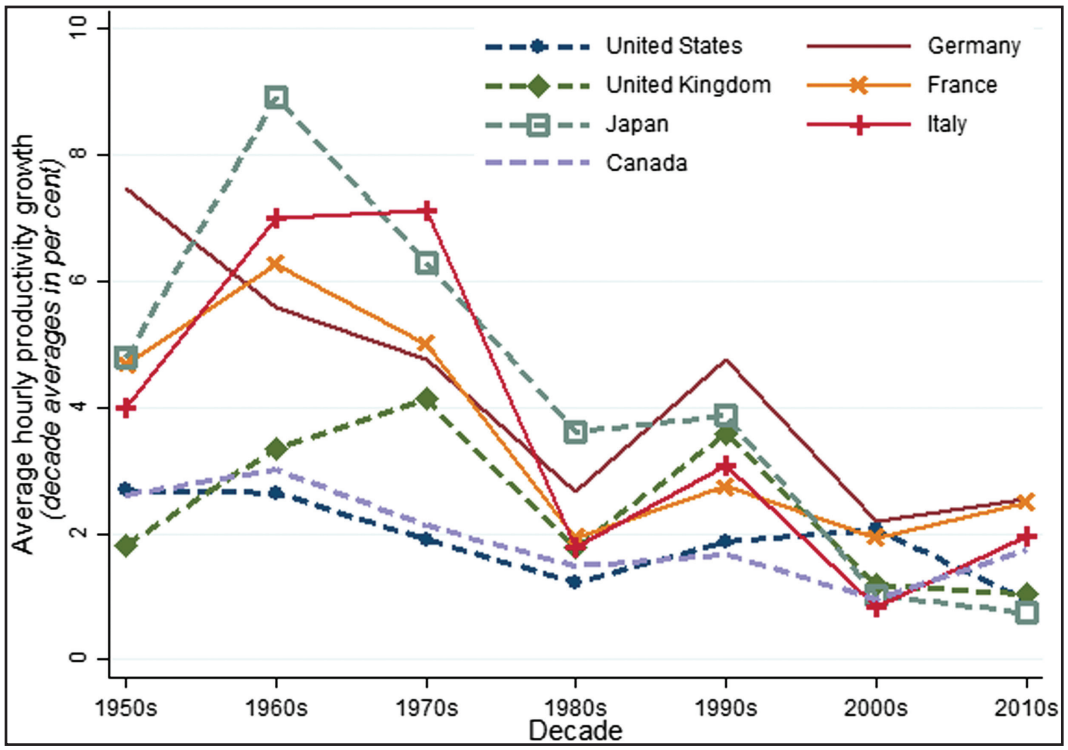

Source: Penn World Tables version 9, available at: https://www.rug.nl/ggdc/productivity/pwt/; authors' calculations.

for labor. According to this scenario, the most affected industry should be manufacturing. This adverse effect might be strengthened by the growing labor quality in developing countries and the ensuing rise in labor costs. The Boston Consulting Group, for instance, reports that wages in China and Mexico increased by 500% and 67% between 2004 and 2014, respectively (Sirkin et al., 2014). This convergence in cost competitiveness is likely to continue in the future, eroding the incentives for producers to move their activities from developed to developing countries.

Offshoring, re-shoring, and robotization are part of a general rethinking of business strategies that have become more complex and based on a wider set of variables than simple cost comparisons (De Backer et al., 2016). On the one hand, the need to face different types of risk and to deal with increased volatility in demand, exchange rates, or commodity prices has shaped outsourcing decisions. These and other issues might have pushed several companies to shore the production back home (e.g., Adidas, General Electric, and Plantronics). On the other hand, the possibility of using cloud-based solutions has reduced the advantage of having low-cost programmers in developing countries. A study of A.T. Kearney has produced projections of job losses in India, Philippines, Poland, and the United States, imputing different automation paces for different outsourced business processes. Their results suggest that countries that have previously benefited from offshoring business processes stand to suffer more job losses than those where this type of job is still onshore (A.T. Kearney Global Services Location Index, 2017).

A shared concern about robotization arose from job polarization, or the fact that middle-skill, middle-income jobs are disappearing to the benefit of job creation both at the high and at the low end of the wage distribution (Autor, 2010; Autor et al., 2003). Such developments toward worsening inequality seem to have eroded the benefits brought from earlier waves of

Figure 3 Job polarization around the world, 2000–2021

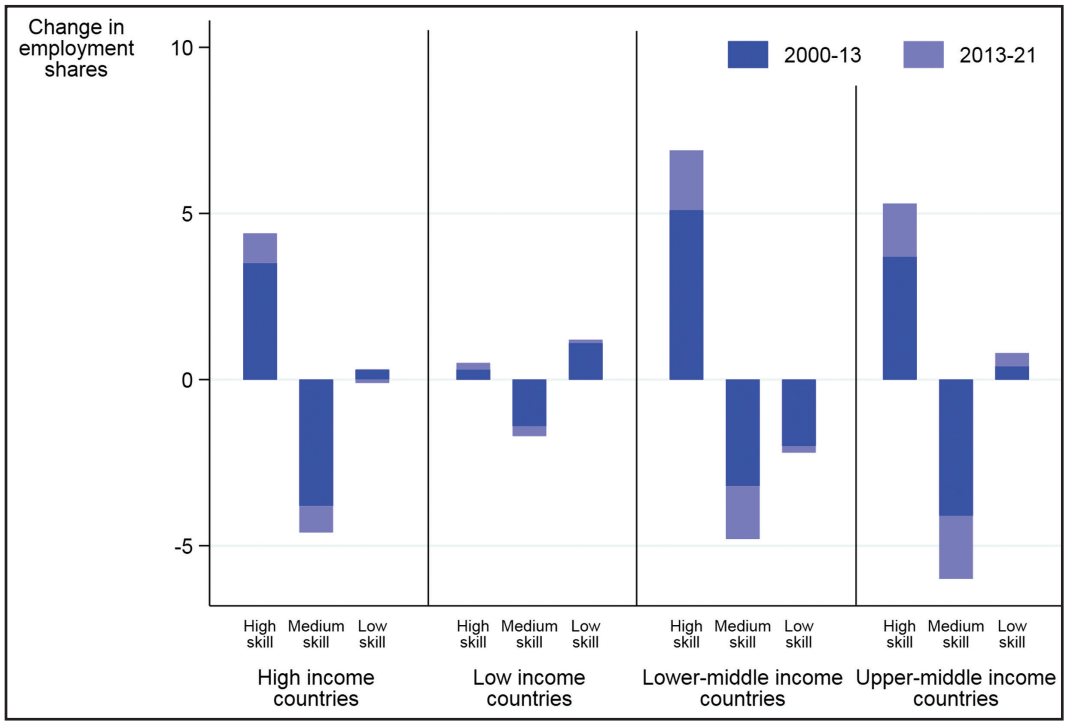

Note: Change in employment shares, in percentage points; forecasts after 2016.
Source: ILO, Trends Econometric Models, Nov. 2016.

productivity increases that lifted the boat for everybody in the long run. Moreover, this change in occupational growth does not seem to affect only advanced economies but represents a widespread phenomenon that is also shared by emerging and developing economies (Figure 3). Indeed, recent evidence suggests that structural change as experienced in advanced economies since the 1950s seems to be characterized by a "hollowing-out" of the middle class, with negative consequences for income inequality and inclusiveness but potentially also for economic development more broadly (Bárány and Siegel, 2018).

The current wave of technological change in the form of AI, therefore, comes at a time when the anticipated benefits from the previous wave have not (yet?) fully been felt and where costs – in the form of higher inequality and lower income growth for the middle class – are becoming manifest. Consequently, concerns are rising that this time, unemployment might actually increase and earnings fall, not least because in periods of stagnating output, increases in labor productivity induced by new technologies necessarily lead to a fall in labor demand. Even if it does not lead to fewer jobs, such shifts could cause working conditions to deteriorate and earnings to fall further behind productivity, as they already have in the past (ILO, 2016). To better understand this, however, we need to look more closely at the linkages between productivity, organization of production, and employment.

3 Jobs, tasks, and the organization of production

When firms automate production, job growth is affected through three channels (Acemoglu and Restrepo, 2017; Chiacchio et al., 2018; Vivarelli, 2014). First, new technologies lead to a direct substitution of jobs and tasks currently performed by workers (the displacement effect);

second, there is a complementary increase in jobs and tasks necessary to use, run, and supervise the new machines (the skill complementarity effect); and third, there is a demand effect both from lower prices and a general increase in disposable income in the economy due to higher productivity (the productivity effect). Typically, these effects do not materialize simultaneously, and the standard narrative runs that unemployment is initially going to rise with automation before falling again when prices and productivity adjust broadly across the economy, often at a much later stage. When distinguishing between different time horizons, these differences in short- vs. long-term effects of productivity growth on unemployment can be indeed discerned in historical trends for the total economy (Semmler and Chen, 2017), even though effects at the industry level might differ and depend on the price elasticity of demand for industrial goods (Bessen, 2017b).

This analysis of how technological change impacts employment is, however, based on three shortcuts. First, it is assumed that when tasks are being substituted by machines entire jobs disappear (almost) immediately. Second, occupational supply is assumed to be inelastic so that a skill-biased change in labor demand induced by technological change will lead to technological unemployment or worsening working conditions (Autor et al., 2006; ILO, 2015); over- or under-qualification does not exist. Finally, the increase in demand that is made possible through higher productivity is supposed to be uniformly distributed across sectors, irrespective of the extent to which these are being automated. In consequence, sectors with higher degrees of automation will experience a relative drop in the share of demand and therefore create less employment, in comparison to those that do not benefit from automation, which again will lead to job polarization and rising income inequality (Bessen, 2018). To understand whether AI will force labor markets through the exact same pattern of adjustment, it is useful to take a closer look at these three assumptions.

3.1 **Changing jobs and tasks**

Jobs are constituted by a set of tasks. If some of these tasks are automatized, job profiles might change by adding new tasks or modifying existing ones instead of suppressing a job entirely. The task description of an administrative assistant over time can demonstrate how similar jobs continue to perform certain tasks that have not (yet) been automatized alongside other, new tasks that either did not exist before or were performed by a different group of workers. Hence, whether or not jobs disappear depends on whether it remains profitable to group certain tasks into specific job profiles and hire workers specifically for these (new) jobs, which is a question more of demand for particular products and services that these jobs are supposed to deliver than of supply of skills to fill the jobs (Acemoglu and Autor, 2011; Bessen, 2017b).

Importantly, cross-country differences exist regarding how jobs are being designed and tasks regrouped into jobs. Ernst and Chentouf (2014) show that tasks have different characteristics regarding their training, supervisory, and production requirements, which are not necessarily aligned. Depending on the importance a company puts on training its workers, supervising them or aligning their workflows, different tasks may be regrouped to jobs from one company to another. Partly, this will depend on country characteristics regarding education and training infrastructure, tax incentives, and social benefit systems (Sengenberger,

1987). Hence, even companies operating in the same industry but in different countries might react to institutional differences with a very different setup of their internal work processes and job profiles, as exemplified by the differences between Apple and Samsung in the way they externalize their production chains. Consequently, whether the automation of tasks will lead to jobs disappearing is as much a technological question as it is an institutional one and cannot be determined a priori by looking at the automation process alone. Recent evidence seems to confirm the importance of institutional factors in determining the outcome of occupational changes, as seemingly similar patterns of job polarization across countries can be driven by different factors (Albertini et al., 2017).

Even when tasks can be automated they might not disappear altogether. Rather than executing a particular task, for instance, an employee might be charged to ensure that the machine is conducting the task properly and to intervene in case of an emergency or error (MGI, 2018a). In the case of air pilots, for instance, the introduction of automatic pilots has not made obsolete their role. Even though on average a pilot only flies a plane for roughly 7 minutes during an entire flight, having a human sitting at the control panel is as essential as before to intervene in extreme situations or sudden disruptions or in technical malfunctions not foreseen by the auto pilot (such as a simultaneous breakdown of both engines).[6] Similarly, it might still require a worker to ensure that machines are properly parameterized and set up, especially when orders change or a new production line needs to be set up. Also, the relative time spent on each individual task might change: thanks to support by AI on diagnosing diseases, doctors, for instance, might spend less time on analyzing symptoms and more time on ensuring a patient's well-being and individual needs. Either way, automation of a task might not necessarily lead to that task no longer requiring human assistance. Rather, the question becomes whether it remains profitable to bundle a set of tasks to a specific job, as well as how quickly a worker can shift within the current job to perform slightly modified tasks or task sets. If that entails requiring new skills that are costly to learn, automation can be expected to lead to inequality within occupations rather than across (Bessen, 2015a).[7]

3.2 **Capital–skill complementarity**

Inequality and joblessness among (low-skilled) workers will also depend on the extent to which machines are complementary to high-skilled labor. The complementarity between skills and machines is not bound by technological factors alone, as the historical account above of different waves of industrial revolutions has demonstrated. Rather, whether or not firms introduce skill-biased technologies depends on whether these are profitable (Acemoglu, 2002). In the 19th century in particular, workers seem to have had comparative advantages over machines in certain repetitive tasks that required high dexterity, for which machines at the time were not yet ready. The relative abundance of unskilled labor at the time made it unprofitable for companies to develop technologies that would allow them to substitute for unskilled labor, as can still be observed in sweatshops around the developing world today. However, as soon as the supply in skilled labor increased and hence relative prices of skilled vs. unskilled labor fell, technologies

6 Source: https://www.nytimes.com/2015/04/07/science/planes-without-pilots.html.

7 Doctors, for instance, might be required to spend more time with their patients, requiring them to develop strong interpersonal skills, something that is not automatically being taught in medical schools.

that made their use profitable began to be developed, leading to the pattern of skill-biased technological change that we can see today (Goldin and Katz, 1998).

With the installation of ever more complex machines, the demand for workers capable of operating and maintaining them rose constantly. Nevertheless, the number of supervisory and skilled workers that these new machines commanded was nowhere near sufficient to create enough jobs to compensate for the loss in demand for the low-skilled workers they were replacing. Hence, capital–skill complementarity became synonymous not only with rising income inequality but also with an increase in technological unemployment to the extent that low-skilled workers were not able to switch occupations or sectors. Most importantly, it was a key explanation of why an increase in the relative supply of the educated workforce did not lead to a fall in the skill premium, that is, the wage difference between high- and low-skilled workers, as one would expect in the absence of such a complementarity. As technological progress gradually reduced the price of capital, investment in new equipment continued and led to a gradual rise in the skill premium.

The extent to which new technologies require the complementary input of skilled labor is, therefore, a main determinant as regards the effect of AI on employment and inequality. Indeed, even modest changes in the degree of complementarity can produce vast differences in labor market outcomes (Berg et al., 2018a; IMF, 2018). To the extent that AI is expected to replace *mental* tasks as explained above, it is, however, not entirely obvious that AI-based innovations might be characterized by strong capital–skill complementarities. Indeed, the entire logic of AI-based systems is to offer expert knowledge to nonspecialists. Whether these systems concern sophisticated medical devices such as activity trackers, agricultural expert systems to guide farmers in selecting and planting the right variety of seed at the right time or sharing platforms for optimizing multimodal transportation, they often require little or no prior knowledge, connect a vast array of users, and provide advice and guidance that help lift productivity, particularly in sectors dominated by low-skilled workers. In construction, for instance, still an area of low productivity that continues to absorb a significant share of low-skilled workers, new computer-based planning systems, for instance, could help to speed up the construction time, cutting waste and optimizing the maintenance cycle of buildings, without changing the skill composition of the sector (MGI, 2017). In other words, part of the promise of AI is that it actually can help lift productivity especially of low-skilled workers, while cutting demand for high- and medium-skilled professionals, quite the opposite of what has been observed in the past.

3.3 **The evolution of demand and the emergence of new tasks**

The rise in productivity that is generated by technological change will help expand incomes and demand. Whether unemployment increases or working conditions worsen will then depend on the types of goods and services this additional demand will be addressed to (Bessen, 2018). Typically, technological change does not progress uniformly across sectors. Hence, the additional income that is generated by automation in one sector might not lead to more demand for that same sector, contributing to a fall in labor demand for that sector. In contrast, if demand for products or services from the automated sector reacts very strongly to changes in price, that is, if demand is highly price elastic, any effects from labor-saving automation might be more than

offset by increases in demand (Bessen, 2018). A recent example is the introduction of automated teller machines (ATMs) in the banking industry starting in the 1970s. Despite the labor-saving nature of the ATM, employment in banking grew continuously as the cost of opening new outlets fell, helping to attract a larger customer base while at the same time shifting tasks among bank employees away from clerk services to sales and counseling (Bessen, 2015b).

Similarly, as demand grows overall, highly price elastic but labor-intensive sectors might benefit, creating additional job opportunities or helping to create new tasks. In the United Kingdom, for instance, demand for recreational and cultural activities has increased by more than 5% points in the consumer basket between 1988 and 2017, in part thanks to the gains made from automation that allowed people to spend less on apparel or food. Similarly, in the United States over a shorter period (1998–2017), spending on health care increased by 2% points in the average consumer basket. Such changes in relative spending patterns toward more labor-intensive sectors can be widely observed and are one of the key factors to explain that technological unemployment has often remained a temporary phenomenon if at all.[8] At the same time, with consumers getting richer, demand for luxury goods and services increases, as can be observed from the steady rise in the numbers of personal coaches and trainers.[9]

3.4 **The impact of AI on jobs and wages**

Taken together, the impact of a large-scale introduction of AI on jobs and wages will depend on three factors: the price elasticity of supply of capital vs. the elasticity of labor, the substitution elasticity between capital and labor, and the direction of technical change induced by AI, that is, whether AI is capital or labor augmenting. The more inelastic the supply of AI, the higher the substitution elasticity between AI and jobs and the more labor-saving AI-based innovations are, the higher will be the extent of technological unemployment and the lower will be any wage gains. Based on the discussion in this section, a nuanced picture arises, in particular as regards the implications of AI for labor markets in developing countries.

First, the elasticity of supply of capital and labor depend to a large extent on how heterogeneous both factors are. The more homogenous a factor input is, the more elastic its supply will be and the less will this factor be in a position to generate high returns.[10] In this sense, skilled labor is less elastic than unskilled one, a key factor behind the wage premium for skills. Similarly, intangibles, such as AI, or robots might not easily be reproducible due to intellectual property rights, data (collection) ownership, or physical limits to consumption of energy and natural resources, which makes the supply of such high-tech capital less elastic. This is likely to be more problematic in advanced economies where overall access to financial markets is well developed and intellectual property rights enforced, leading to a low relative price of traditional capital. In developing countries, on the other hand, the capital price of AI relative to

8 Data on consumer basket spending items are taken from ILO statistics.

9 Absolute numbers are small, though, despite a global growth rate of around 12% between 2011 and 2016. Currently, an estimated 53,300 people are classified as personal coaches and roughly 128,300 people have part of their tasks related to coaching (ICF, 2016).

10 A. Marshall uses the concept of "quasi-rents" to describe excess returns over and above the marginal product that will erode over time as factor supply adjusts. In this regard, the less elastic factor commands a higher quasi-rent and will benefit more from the increase in productivity. In modern, search theoretical approaches to the labor market, quasi-rents are linked to the degree of specificity that is determined by both search/transaction costs and the value of the outside option (Marshall, 1890).

traditional capital is likely to be lower, given more restricted access to capital and higher risk premia overall as regards investment. Investment in AI might, therefore, be relatively more elastic given the generally higher profits to be made in such an environment. At the same time, developing countries still have large supply of unskilled labor, which prevents wages from rising (faster) but which also reduces incentives to invest in AI technology. Only when the supply of (unskilled) labor slows down, the incentive for a shift toward automation will become stronger, as it can be currently observed in China and other emerging countries (see Carbonero et al., 2018, who document the rapid rise in robotization in some of these countries).[11]

Second, a high elasticity of substitution between capital and labor leads to a reduction in labor demand with the introduction of new technologies. Previous waves of high-tech innovations came with a strong complementarity between capital and skilled labor, leading to increases in wage premia and job polarization. As we argued before, however, with AI the degree of complementarity between capital and skilled labor might actually be lower as AI has the potential to increase the productivity of low-skilled labor. At the same time, some AI-based applications are replacing tasks carried out by medium- and high-skilled workers, which might lead to a reduction in wage premia for skilled workers, thereby weakening pressure for job polarization. In other words, lacking skills might actually not be a barrier against the use of AI and hence stimulate the demand for this type of labor. On the other hand, skilled workers might no longer benefit from a complementary relation with capital, in particular if their skills do not match with the requirements for the development and implementation of new AI applications (Acemoglu and Restreppo, 2018).

Finally, to the extent that investment in AI is capital or factor augmenting it will increase capital productivity or scale up production without displacing labor. In this case, the productivity effect is stronger and leads to more jobs and higher wages, albeit the impact on the wage premium for skill labor is unclear. In the case of labor-saving technical change induced by AI, however, the situation is more complex as labor is replaced and the overall impact on labor markets depends on the size of the productivity effect and the extent to which induced demand is big enough to compensate for displaced labor. As discussed previously, the impact of labor-saving technological change on labor demand will also depend on the price elasticity for the goods and services that are being automated: to the extent that automation happens in (services) sectors with large unmet demand, price elasticity might be high and a reduced price thanks to automation will lead to a strong increase in demand that compensates for the substitution effect. Moreover, in the next section, we discuss that many applications of AI are capital and factor augmenting rather than labor saving, for instance, when they improve the matching process on different (labor and product) markets and enhance the productivity of installed capital (for instance, in the energy sector).

Considering these three factors leads to a more optimistic outlook as regards the impact of AI on jobs and wages, in particular when looking at its potential to support the catching up process in developing countries. The extent to which AI supports labor demand and wage growth will, however, depend on the concrete applications that are currently being developed. Moreover, the distributional consequences of AI are linked to broader considerations about the

11 The point at which real wages start to accelerate in the process of a country's economic development is also known as the "Lewis turning point," at which the supply of low-skilled labor slows down or declines, for instance due to slower population growth, lower internal migration from rural areas, or a general increase in the level of education and skills (for the Chinese experience, see Zhang et al., 2011).

implications of the rise of intangibles – which AI capital belongs to – and competitive forces on product markets. This is what we will turn to next.

4 What is different about AI?

Can we expect AI to have labor market effects similar to previous waves of automation such as those resulting from robotization? Many observers believe, indeed, that AI – given its focus on mental rather than physical tasks – has the potential to become another "general purpose technology" with a wide range of applications in various sectors and occupations (e.g., Furman and Seamans, 2018). This could mean that the results we have found so far, given the fact that they are based on robotization of only a few sectors, could generate even more significant (negative) employment effects when AI affects a far larger set of industries and occupations. However, as we have argued at the beginning of this paper, not all the insights that studies on robotization have generated might carry over where AI-based technologies are being developed and adopted more widely. Most notably, whether AI-based technologies are characterized by the same degree of capital–skill complementarity as robots is not entirely obvious. In this section, we look more closely at the specific applications that look feasible from a current perspective, and the potential labor market implications, making use of the discussion in the previous sections.

4.1 Specific characteristics of AI

As discussed in Section 1, the development of AI has benefited from three interrelated trends: the availability of large (unstructured) databases, the explosion of computing power, and the rise in venture capital to finance innovative, technological projects. These have allowed the rapid development of new applications in areas where humans were thought to have a particular advantage: making predictions and taking decisions regarding routine yet nonmechanical tasks. Typically, these types of tasks were mainly found in the services sectors, which employ – even in emerging economies – more than half and sometimes up to 70% of the workforce. Three main groups of tasks have become the focus of AI applications, in particular:

- *Matching tasks*: The most prominent group of tasks concerns all those jobs that consisted in matching supply and demand, especially on markets with a heterogeneous product and services structure. Whether ride-hailing services (Uber, Lyft, Didi Xiuching), hotel and accommodation services (AirBnB, Ebookers, Booking.com), retail (Amazon), or human resource management (LinkedIn) among others, machines have proved to be significantly faster and more efficient in identifying matches in these markets. This, in turn, helps companies to cut costs on finding customers or suppliers and offering less expensive solutions to their growing customer base, often, however, at the cost of worsening working conditions of their suppliers and their employees. In particular in the gig economy, where demand for micro-tasks such as image classification or survey responses is matched with workers available for short-term, on-demand tasks, working conditions are often below (national) minimum conditions (Berg et al., 2018b). An additional concern arises where privacy rights are not or insufficiently being protected, leaving employers in a strong position to (further) undermine worker rights and working conditions (De Stefano, 2018).

- *Classification tasks*: Early applications of AI concentrated on image and text recognition techniques, especially facial recognition, partly in relation to the increase in surveillance cameras and techniques. In the meantime, however, an explosion of applications has taken place in this area, including in medical applications (X-ray image diagnosing), legal services (reading and classifying legal documents), accounting and auditing (analyzing balance sheets, fraud detection), recruitment (screening applicants), and potentially threatening the jobs of a significant number of well-paid workers in the services industry. Yet, it also promises to enhance significantly the productivity of the most productive workers in these industries even further: automatic text generation software allows journalists and editors to concentrate on those key, high-valued added papers that attract a large customer base to their employers. Similarly, automatized research designs help scientists to focus on the most promising areas of their experiments (for instance, in the development of new drugs) while allowing the computer to discard all those research avenues that are likely to fail (Cockburn et al., 2018). The democratization of expert knowledge that these AI applications bring, however, also runs the risk of expert deskilling and abuse, for instance, in the case of facial recognition, which has recently led industry leaders to call for a careful regulation of these technologies.[12]
- *Process management tasks*: A final set of applications concerns a combination of the two previous sets of tasks, identifying patterns and bringing different suppliers and customers together along a supply chain (Culey, 2012). This type of complex network management also arises in the management of electric grids and complex infrastructure and building projects, including the maintenance of finished projects (through the Internet of things, IoT) or multimodal transportation solutions to curb inner-city traffic. In combination with decentralized tracking and certification schemes (Blockchain), it includes the implementation of expert systems across supply chains, allowing upstream producers to integrate diversified supply chains through better information about product quality, certification schemes, and market conditions. These types of expert and complex management systems are of particular relevance in developing and emerging countries, helping local producers to gain access to a wider set of expertise on production conditions, supply chains, or simple learning tools.[13] It is this latter group of tasks that currently bears no resemblance to what robots used to automate in the past. Rather, these new AI-based innovations constitute a new group of tasks that either cannot be properly carried out by humans due to their complexity or have been too expensive to be performed by human workers, even in combination with traditional technologies (Benhamou and Janin, 2018).[14]

12 See https://blogs.microsoft.com/on-the-issues/2018/07/13/facial-recognition-technology-the-need-for-public-regulation-and-corporate-responsibility/.

13 The potential of applications of AI to developing countries has already been recognized by major tech companies. Google recently announced that it was to open its first African AI research lab in Accra (Ghana) to develop tools specifically designed for local market conditions; see https://www.blog.google/topics/google-africa/google-ai-ghana/

14 One of the insights from the earlier endogenous growth literature was indeed that new types of goods and services become available only once they are sufficiently profitable to be carried out. In other words, labor demand for the production of certain products or services is essentially zero at any point in time to the extent that current technologies do not allow them to be carried out profitably. By one account, AI has added up to 7% of GDP in the United States due to these additional services that were hitherto not accessible to humans (Cohen, 2018).

Without stretching the task-based methodology discussed previously too much, these three fields of applications of AI can be categorized as (a) task substitution; (b) task complementarity; and (c) task expansion. In the case of matching applications, existing tasks are being taken over, often in a more efficient way, through algorithms that allow the matching of supply and demand more rapidly and more precisely. In the case of classification tasks, AI-based applications help workers involved in such tasks to concentrate on those that require specific attention while leaving the more routine, repetitive tasks to a computer. Finally, as regards process management tasks, here AI-based applications often carry out tasks for which no human workforce was available to begin with, precisely because of the complexity of the tasks; in this case, the computer essentially expands the number of tasks that are being carried out in an economy, thereby enhancing total factor productivity regardless of whether production is based mainly on skilled or unskilled labor. A priori, therefore, it is not possible to determine whether the development and diffusion of AI-based applications will contribute to widespread job destruction or to an increase in inequality. The effects of AI will depend on the relative importance of these three different areas of applications of AI. In particular, they will depend on the direction that technological change will take in the future, under the impression of policies, tax incentives and public and private investment in technological research (Mazzucato, 2013). In other words, the extent to which AI will lead to a recomposition of tasks and jobs will partly depend on the particular technology and innovation policies in place to orient the technological progress in socially desired ways. We will get back to this point in our final section on policy options.

4.2 **The economic and social implications of large-scale applications of AI**

The large-scale application of AI might yet generate additional economic and social implications, irrespective of whether these applications are substitutes, complements, or extensions of existing tasks. These implications have to do with the particular nature of AI: AI is digital in nature and therefore non-rivalrous, similar to other digital products and services, that is, digital services can be used by more than one person without affecting each other. Moreover, AI aims at providing individual solutions to economic problems, not only allowing for a more enhanced product and service diversification than ever seen before but also for much finer price discrimination than on existing markets. Such price discrimination is, however, a double-edged sword, as the additional opportunities it might provide for some have to be compared against the proliferation of preexisting biases this might entail. Nevertheless, and related, the use of AI in helping to reduce matching frictions – irrespective of its task substitution nature – also creates more opportunities for market interconnection and exchange. Finally, AI systems by their very nature represent embodied technological change, with specific implications for the skill-biased nature of this form of economic progress. Let us look at these issues in more detail.

First, digital technologies that are characterized by non-rivalry in the use of their products and services often provide cumulative advantages to those entering first a particular market (segment). Once fixed costs for the development of new digital services are being deployed, a growing market can be served (almost) at zero marginal costs, with economies of scale significantly larger than during previous waves of technological change based on automation

of mechanical tasks (Moretti, 2012). This gives rise to superstar firms where few companies dominate and occupy a privileged, highly profitable position, potentially limiting competitive pressure by erecting barriers to entry (Rosen, 1981; Autor et al., 2017a, 2017b). Second-movers often face uphill battles to enter the market or have to focus on small market niches with less profitable opportunities, producing large inequalities between individuals and between firms. Korinek and Ng (2017) argue that recent technological changes have transformed an increasing number of sectors in the economy into the so-called "superstars sectors," in which a small number of entrepreneurs or professionals concentrate the demand of a large range of consumers. Examples include the high-tech sector, sports, the music industry, management, finance, etc. Importantly, these superstar dynamics are not limited to firms producing digital goods and services, but increasingly include those using them, thereby affecting a potentially much larger group of sectors and occupations. As a result, superstar firms and employees concentrate enormous rewards in a wide range of activities, widening the gap with the rest of the economy and reducing the share of income to labor (Autor et al., 2017a).

The superstar dynamic is further reinforced through business practices that enhance the first-mover advantage. Indeed, some companies are adopting data-driven business models and strategies to obtain a competitive "data advantage" over rivals. Data-driven mergers (e.g., Facebook's acquisition of WhatsApp) are increasing the risk of abuses by dominant tech firms. Data-driven exclusionary practices and mergers raise significant implications not only for privacy and consumer protection, but also for competition law. Due to network effects, data-driven mergers may increase entry barriers and enable some big firms to become bigger until they dominate the whole industry (Stucke and Grunes, 2016). In this light, some commentators within the antitrust community are raising concerns about the potential harm of data-driven mergers and abuse by dominant companies built on data. The Organisation for Economic Co-operation and Development (OECD) recently warned that data-driven markets can lead to a "winner takes all" result (OECD, 2015a). These network-driven market concentrations are likely to grow larger with AI, which is very much based on large, centrally available databases.

A second source of change comes from the fact that AI-based systems allow for a much finer discrimination between different customer groups. Indeed, market segmentation and differential pricing is nothing new and has been practiced for some time. However, AI allows firms to predict individual costumer behavior and price sensitivity in much more detail. Based on previous consumer and search patterns, for instance, on online shopping platforms or as revealed by credit card transactions, suppliers can essentially charge individual prices or suggest individualized price–service quality combinations that allow them to reap a much larger part of the consumer surplus than in the past. Such so-called third-degree price discrimination has not yet been a matter of active research in relation to AI but some insights from previous research allow a couple of conclusions to be drawn (see Tirole, 1988; Gifford and Kudrle, 2010).[15] With this form of price discrimination, producers offer (groups of) consumers the same type of product or service at different prices, based on the relative willingness of consumers to pay for these products. A typical example consists of internationally traded goods, such as pharmaceuticals, that are priced differently depending on a country's consumer characteristics,

15 Arguably, AI-based price discrimination could be considered as first-degree price discrimination, allowing full extraction of consumer rents by producers. This, however, would require perfect prediction of a consumer's willingness to pay, something that runs counter to the underlying principle of AI systems as stochastic prediction machines.

which may depend on differences in regulation and taxation. One general conclusion from this research is that welfare increases if and only if the total output produced by serving different market segments at different prices exceeds the output in a situation where all consumers pay the same price. It turns out that this is the case under fairly general conditions, but it also implies a shift of (part of) the consumer rent to producers, thereby worsening any prior trends toward higher levels of inequality (Varian, 1985).

Recent developments have aimed at applying this to human resource management as well. Indeed, the area of what has become known as "Human Resources (HR) analytics" aims exactly at this type of price discrimination to attract workers to companies, differentiating between categories of employees in terms of working conditions, wages, fringe benefits, or responsibilities. A particular concern with this type of discrimination of working conditions arises from the fact that reservation wages of different groups of otherwise similar jobseekers may be caused by past discriminations observed in the labor market. Women or ethnic minorities, for instance, might be ready to accept lower wage offers, as they were experiencing higher entry barriers in the past. An automated recruitment system based on analyzing historic data would replicate this type of bias, thereby reinforcing preexisting discrimination (Ponce Del Castillo, 2018). Hence, even though price discrimination might, in general, allow expansion of the number of available jobs, it is suboptimal in cases where differences in willingness to pay (or to accept job offers) depend on previous discriminatory practices. So far, however, it seems that people continue to hold favorable views of algorithmic decision-makers vis-à-vis humans, suggesting that even though algorithms come with their own biases, these might be (seen as) less harmful than those perpetrated by humans (Logg et al., 2018).

At the same time, however – and this is a third area of economy-wide applications of AI-based systems – matching frictions on labor markets can be substantially reduced when automated systems allow a significantly larger pool of applicants to be processed. Indeed, mobility of workers, whether across occupations, sectors or locations seem to have declined in recent decades (Bunker, 2016; Danninger, 2016; Molloy et al., 2014). Part of this fall in labor mobility has to do with regulatory barriers such as occupational licensing or barriers to geographic mobility. But a significant part relates to informational frictions and difficulties for employers in properly identifying competencies from past experiences or education. Similar to applications in the area of HR analytics discussed above, AI-driven matching systems are helping to identify the appropriate mix of firm internally and externally available competences to bring them together for specific projects such as the development of new products or services. Indeed, AI has already started to shift the boundaries of the firm in favor of more and more services being insourced from external (labor) markets, such as through micro-tasks available through gig platforms (Berg et al., 2018b). Job search platforms such as Monster.com or LinkedIn are already offering detailed models of job vacancies and available candidates to help recruitment managers and applicants in matching job requirements with candidates' (self-declared) competences. The benefit of using AI in this area comes not only from the larger pool of applicants and vacancies that can be matched against each other (thereby enhancing labor market fluidity). It also lies with the improved identification of competences based on self-declaration and historic professional experiences that might be difficult for an individual recruitment manager to properly discern. So far, these systems still seem to be far from perfect and riddled with biases, as anyone who has used them can confirm. Nevertheless, the expected

efficiency gains promise to be large: according to MGI (2015), for instance, enhanced matching efficiency thanks to such online job platforms could yield an additional 72 million jobs worldwide and spur global GDP by 2% within the next decade. Notwithstanding, these efficiency gains must be matched against a likely increase in employment volatility and job insecurity, especially when such newly created jobs are of only temporary nature.

A final, economy-wide implication of AI concerns the fact that technological change driven by AI is embodied in new and often cheap equipment, accessible to a wide range of users.[16] Its digital nature and the fact that many AI-based expert systems can be run from currently available mobile phones have contributed to its significant diffusion, including among users in emerging and developing countries. The particularly steep fall in capital prices that is being fueled by AI is likely to help boost productivity especially in those regions and parts of the world where lack of finance and other barriers have prevented the implementation and diffusion of existing technologies. As discussed above, expert systems are currently being developed to help, for instance, smallholder farmers to get better information on what, when and how to seed to improve the agricultural yield. In particular in certain semiarid regions in Africa, precise advice on meteorological conditions in combination with proper farming and irrigation techniques has been shown to yield substantial potential for productivity gains through water savings and more appropriate seeds.[17] Given that today more than one-third of all workers worldwide still work in the agricultural sector, such productivity increases promise to alter significantly the development potential and income opportunities, including among low-income countries. Similarly, using AI-based matching and supply chain systems holds the potential to cut down on logistics and transportation costs, an issue particularly relevant for producers in developing countries that often lack access to large distribution networks.[18] Finally, the delivery and implementation of public policies often depends on timely and precise information about areas in need of intervention. AI-based expert systems have been shown to help policymakers, in particular in countries with limited fiscal resources, in better managing their interventions, delivering better, more granular information, and allowing an improved coordination of various actors necessary to, for instance, deploy medical care or emergency interventions.[19]

5 Policies

The previous sections have demonstrated the wide and varied job-specific and economy-wide implications of AI that result from its general purpose nature. AI's potential to generate major productivity enhancements, in particular in sectors and countries that so far have not benefited from significant structural change, has to be matched against the risk of worsening gaps in income inequality as first-mover advantage looms large and can easily be reaped. The following section discusses some of the policy implications that this assessment warrants. Specifically, it

16 Economists distinguish between embodied and disembodied technological change. The former relates to all those forms of innovation that are being implemented through investment in new tools, machines, and equipment. The latter arises from innovations in the way existing labor and capital is being organized, for instance through organizational innovations or innovations in infrastructure and regulation that help to make more efficient use of existing technologies.

17 See, for instance, the Tunisian start-up iFarming, http://www.jeuneafrique.com/501309/economie/start-up-de-la-semaine-ifarming-future-licorne-tunisienne-de-lirrigation-en-temps-reel/.

18 https://medium.com/@KodiakRating/6-applications-of-artificial-intelligence-for-your-supply-chain-b82e1e7400c8.

19 https://channels.theinnovationenterprise.com/articles/ai-in-developing-countries.

will focus on four areas of policy interventions: (1) supporting the adjustment of the workforce to be able to transit to jobs and tasks in which workers continue to benefit from a comparative advantage, while being able to make use of the new technologies; (2) guaranteeing an equal playing field between firms by maintaining a competitive environment and preventing individual companies from reaching market dominance, a tendency that has already worsened inequality and hampered productivity growth; (3) reinforcing existing tax and social protection systems in order to mitigate both the impact of the ongoing transformation of the world of work as well as the deepening of income inequalities; and (4) enhancing international cooperation and social dialogue to broadly share technological rents.

5.1 Skills and occupational mobility

The current education systems need to be examined given the arrival of the AI-based wave of technological change. Its current setup as a young age, once-and-for-all type system of skill provision is no longer sufficient when it comes to retraining workers who expect to have an increasingly lengthy work career. Most current proposals, however, start from the premise that what is required is a general uplifting in technical skills for workers to be able to cope with the coming changes. The previous discussion has argued that this is not necessarily the case beyond the capacity to use these new technologies. Importantly, even if the expected increase in the demand for technological skills materializes as currently predicted, social and emotional skills remain the dominant driver for total hours worked, at least in advanced economies, according to a recent study by McKinsey Global Institute (Fig. 4).

Figure 4 Predicted shifts in skill sets, United States vs. Western Europe, 2016–2030

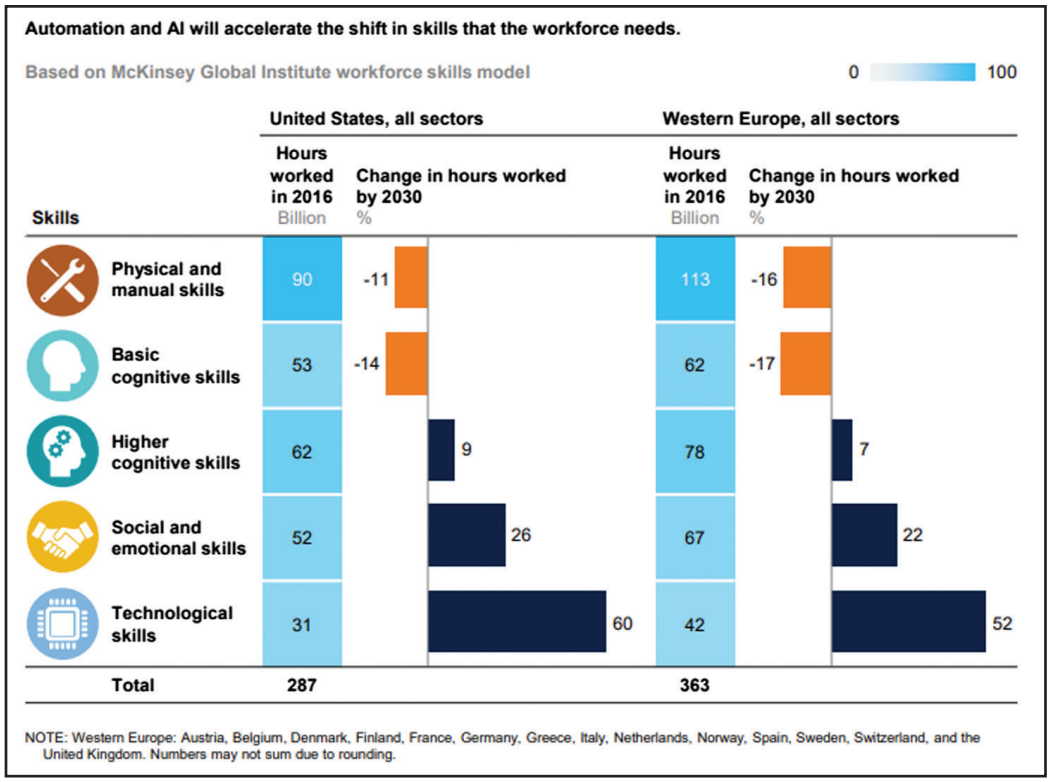

Source: MGI, 2018b.

This dovetails well with the general considerations developed previously about the generic nature of AI-driven technological change. Indeed, technological skills will mainly be asked for in areas where new digital products and services are being developed, which by the nature of this digital industry will remain relatively limited. However, in the areas of application and use of these technologies, new opportunities emerge. In this regard, a certain generic understanding of the availability and use cases of new technologies will be necessary as a broad skill, much as reading and basic mathematics skills are considered to be required for today's low-skilled workforce. However, in an age where there are more mobile subscriptions than actual users and a penetration rate of smartphones of more than 60% of the total population in most advanced economies, many users are already exposed to new technologies and dispose of basic numeric skills.[20] As routine tasks such as verification, compliance, and system processing are increasingly being taken over by machines, human work will shift toward sales, market development, and consulting/coaching, all of which are tasks that require strong social, empathic, and interpersonal competences rather than relying exclusively on technical skills. The latter will still be necessary, but mostly in order for workers to *use* rather than to *develop* new technologies.

These are not new competences, and social and emotional skills have been emphasized by employers in the past. Indeed, an increasing need for social skills has already been observed over the past decades (Deming, 2017). However, current education systems with their strong focus on providing technical skills will need to integrate competence development in this area to a larger extent than in the past. At the same time, this shift in the skill basis also holds the promise that even those people who might find it challenging to access highly technical skills will have a higher chance to integrate into the labor market successfully, provided that they hold the right social and interpersonal skills. In this regard, AI-driven technical change will not necessarily be as skill biased as the previous wave of digital technologies. In particular, in those countries where only few people possess the right technical skills to contribute to the development of AI applications, users of these new tools can expect to enter the labor market successfully even with a diverse and nontechnical skill set.

This is especially promising for currently low-income countries that often do not possess the resources to set up education systems with a similar scope and breadth as more advanced economies. In these countries, AI-based tools can play a particularly productive role in overcoming educational challenges, as they allow local consumer behavior and production characteristics to be sourced to provide tailor-made solutions, for instance for smallholder farmers. Indeed, whereas previous generations of expert systems were often based on hardwired expertise gathered in different countries and contexts, the learning capacity of AI tools makes them particularly amenable to be deployed in a variety of situations without much prior knowledge about local circumstances. Local users of these technologies, therefore, are not required to know much about the underlying technology, nor need they provide sophisticated input into such devices. Rather, their day-to-day usages will allow AI-based tools to generate advice based on overall best practices in combination with local circumstances. This creates low entry barriers for the diffusion of these new technologies and allows training and education to be focused on basic numeric and literacy skills. Hence, even though developing countries might find it

20 See http://resources.newzoo.com/global-mobile-market-report.

challenging to upgrade their education systems quickly and thoroughly enough to expect to be able to produce AI applications, even with limited resources they might expect to be able to use these applications on a broader scale, with large benefits for their growth potential.

A final point concerns occupational and geographic mobility. As new applications will emerge in yet unknown areas of the labor market or new locations, maintaining fluidity between occupations and geographical areas remains important. In this sense, activation and education systems need to account for flexibility both across and within occupations over a lifetime and between locations. Younger generations currently entering the labor market can and do expect to work until their mid- to late 60s, including in emerging economies.[21] Education systems that provide skills only at a young age are unlikely to fit the purpose of an ageing society with (fast) technological change (Agrawal et al., 2018b). Several attempts have already been made to promote (incentives for) lifelong learning, but opportunity costs are typically very high for workers in their prime working years, and skill provision for those on a job search often focuses on a speedy return to employment rather than a more long-term sustainable solution to any shortcomings in skills. Activation systems more broadly need to integrate the perspective of employability over the life span with a focus on competence development that can be used across a range of locations and possibly countries.

In this regard, education systems will need to focus increasingly on competences rather than on skills and promote the certification and portability of these competences. Partly, this will require a widening of the currently narrow occupational licensing that continues to hamper successful labor market integration, even in the absence of AI-based technological change.[22] Moreover, international coordination on a broad set of competences will be required to allow for more labor mobility and better international comparability of those competences, which should help workers more easily to find employment opportunities in new occupations, sectors or locations. Recent initiatives to develop "skills passports" allow to document and certify competences acquired on the job. These could be extended toward a broader, potentially mandatory industry- or nation-wide scheme that helps workers over their working life to assess and identify both their current competences and possible gaps in light of a labor market transition.[23]

5.2 **Ensuring a level playing field among firms**

Besides ensuring a properly prepared workforce, policymakers also face the challenge of maintaining a dynamic labor demand. As discussed above, the digital nature of AI creates significant and persistent first-mover advantages that deepens the gap between early adopters at the technological frontier and the remaining firms. As a consequence, productivity differentials have widened across all OECD countries and firm-level concentration has increased globally,

21 Although not strictly the subject of this paper, the issue of population ageing cannot be abstracted when discussing labor supply and incentives for (higher) education. High levels of educational investment pay off when people grow older. At the same time, given (fast) technological change, educational obsolescence requires constant (and more important) renewal of competences and skills as average retirement ages recede. Despite a long-lasting recognition of the importance of lifelong learning, so far very little has been undertaken to allow workers to benefit from a continuous upgrading of their skills.

22 For a summary of the effects of occupational licensing on the foreign born in Germany, see Runst, 2018; for an overview of their effects in the United States, see https://www.brookings.edu/research/occupational-licensing-and-the-american-worker/.

23 See, for instance, the initiative at the European level, the Europass: https://europass.cedefop.europa.eu/documents/european-skills-passport.

with potentially pernicious effects on productivity growth and job creation (Andrews et al., 2016; Autor et al., 2017a, 2017b). Large productivity differentials between firms have, indeed, been shown in the past to constitute a barrier for wider diffusion of technological progress and innovation among lagging firms, a pervasive problem in countries with large informal economies (Aghion et al., 2005; Boone, 2001). The danger is that with the early adoption of AI-based technologies in leading companies, the productivity differential is set to widen, leading to a rise in market concentration and a push toward "informalization" of those companies that are falling further and further behind the productivity frontier, with consequences for wage growth and working conditions. In addition, the concentration of profit and wealth among a few, large companies creates the risk of regulatory capture by the rich, with adverse consequences for open markets, innovation diffusion, enforcement of (labor) regulation, and country's capacity to collect taxes (see Naudé and Nagler, 2015).

Establishing and maintaining a competitive environment for AI to benefit the economy more broadly can be achieved through three different policy measures.

5.2.1 **Investing in digital infrastructure to share the benefits of AI more broadly**

Investing in digital infrastructure is a key measure to ensure that companies across a broad spectrum of sectors and locations can successfully compete. For emerging economies, this creates new opportunities, as in the absence of a legacy infrastructure (e.g., in high-speed fiber-optic Internet) new public infrastructure can be deployed without interference from an incumbent, thereby helping to create a level playing field. Certain successful non-AI innovations in electronic payment systems (M-pesa in Kenya) or electric vehicle development (China) can testify to the success of such a strategy. However, even in the presence of an incumbent, policymakers need to ensure that the latest infrastructure with high scalability is being deployed to allow companies to take full advantage of the new technologies. Without such an (public) investment in digital infrastructure, the applications and deployment of AI will remain limited, in particular in developing countries where such infrastructure is significantly lacking.

5.2.2 **Providing basic AI tools in the form of open source to enhance access to AI for all**

AI is first and foremost a set of (statistical) methods that need to be implemented in a concrete business case. Often, however, the initial step to explore and evaluate opportunities that arise from AI might not be fully anticipated by market players, especially if they are not operating at the technological frontier. This might be a particular problem in developing and emerging countries. The tech industry itself has indeed recognized this problem, and leaders in this field have offered their expertise and patents on a nonprofit basis to develop new applications with external partners.[24] Such platforms can help prevent private companies from occupying niche areas with large social externalities (such as in the development of new medication). Most importantly, it could allow public technological institutions to codevelop new applications that would help start-ups or other market entrants to compete successfully with incumbents.[25]

24 See, for instance, Elon Musk's open AI initiative: https://www.openai.com/.

25 Germany's public system of Fraunhofer Institutions that collaborate with private partners in the development of industrial applications are an example for such collaboration that predates the current wave of AI. Similarly, the Netherlands has a system of private–public partnership in scientific research and technological development conducted by the Netherlands Organisation for Applied Scientific Research (Toegepast Natuurwetenschappelijk Onderzoek, TNO).

Keeping access to official statistics and basic AI functions as a public good will be essential in maintaining a competitive environment and preventing further industry concentration. Governments could, for instance, pursue an Open Data Policy that would guarantee free access to official statistics, including (anonymized) access to large micro datasets, which currently often require significant amount of time and financial resources to work with. Such a policy could be complemented by setting up public (research) institutions that help codevelop new algorithms for Big Data analysis under the requirement that these algorithms remain open access as well.

5.2.3 Adjusting antitrust policies to prevent first-movers from establishing market-dominant positions

A final set of policy measures consists in adjusting antitrust legislation and intellectual property rights to the particular challenges posed by the digital economy and AI in particular. In this regard, the question of how to properly account for, price and tax the data input that is essential for developing and training new AI algorithms constitutes a key issue. In addition, the legal framework needs to be extended to not only grant ownership to data but also to the predictions generated from these data. This will have direct implications for the sharing of technological rents.[26]

Currently, intellectual property around AI is governed by different regulations and laws. Data (collections) are protected by copyright laws, whereas AI algorithms fall under the premise of patents, which are characterized by stricter time limits (and hence potentially weaker protection). The output of AI tools (for instance, creative works) are, so far, not protected. Similarly, individual data are not being protected (but rather are considered confidential), including against false information (Scassa, 2018). This is a particular challenge when workers, customers, or debtors are shunned from market opportunities because of false information recorded in the databases on which AI algorithms base their assessment. Accounts of credit scoring systems shunning potential debtors from financial services because of erroneous information regarding birthdates or names are well known. With more and more matching taking place through algorithmic processes, these problems are likely to multiply as market participants do not have the possibility to verify and potentially contest the data recorded about their digital profiles (avatars). Extending the existing copyright framework to individual data may, however, be too strict as it would also prevent the development of efficiency-enhancing applications. Rather, a more balanced protection of ownership across different categories (data, data collection, algorithms, services) seems to be necessary to balance better the need for privacy and data truthfulness with business interest to innovate and develop new products and services.

In this regard, switching costs between networks prove to be a particularly challenging issue that locks customers to specific service providers (Stucke and Grunes, 2016, ch. 10). For instance, signing up and matching with potential employers on gig work platforms entail substantial costs for workers (Berg et al., 2018b). If the time and energy invested to create and update a profile on one particular platform cannot be transferred to another one (for instance, because the platform does not allow the download and transport of the entire network tree of a particular worker), this creates a significant position of dominance for the platform provider and distorts the terms of trade in its favor. Given that the value of such a network and the

26 https://www.techworld.com/data/ip-rights-for-ai-who-owns-copyright-on-content-created-by-machines-3671082/.

precision of the AI matching algorithm depend on the number of network members, the first to open a platform and attract members enjoys a substantial advantage over possible competitors and can pocket a significant share of the consumer surplus.

Another issue arises from the ownership of the products and services produced by AI. As mentioned, copyright law currently does not protect work produced by nonhumans (the famous monkey who reproduces a work by Shakespeare by accident). Therefore, AI-based creativity, for instance, in plastic arts or music, is currently not being covered by copyright law. At the same time, individual software code used to produce these works is highly protected, limiting the replication, diffusion, and use of this software in a different, potentially more productive, use. Again, the market barriers that this creates are substantial, besides distorting the incentives for developing AI rather than benefiting from its outcomes. Specifically, this creates significant adoption barriers for applications of AI in low-income and developing countries, despite the potentially large benefits they could confer in these countries. In both cases, therefore – the portability of network information and the protection of copyrights – legislation is bound to evolve to recognize the new reality and to weigh the benefits of open competition against the challenge of worsening (pre)existing inequalities. Open source projects and the development of Creative Commons as an alternative to traditional copyrights can be seen as first steps in this direction.

Taken together, this discussion suggests to move the current intellectual property rights system away from the (strict) protection of upstream data input toward patents and copyrights of downstream products and services. This would help strengthen competition at the insourcing of new data, while giving data providers strong incentives and the means to ensure data truthfulness. Lower switching costs and stronger competition around algorithm development would erode current monopoly rents and while enforcing at the end consumer of digital products and services would shift business models back toward more traditional pricing models. Potentially, this will require the development of data industry standards in order to allow a smooth interoperability of different data systems.

5.3 **Social protection and taxation to tackle inequality and job polarization**

Providing support for those in transition and ensuring social cohesion through a reduction in income inequality remains a key challenge given the AI-based wave of technological change. In this regard, tax–benefit systems play a key role in helping workers to cope with transitions to new opportunities in different occupations, sectors, or locations. This seems particularly important in light of the reduction in labor mobility discussed above, that entails significant adverse consequences for the possibilities of workers to benefit from new opportunities. Besides limits to occupational mobility related to lack of skills or industry concentration, the tax–benefit system and other institutional barriers are hampering mobility as well. Portability of benefits, including within the same jurisdiction, is not always guaranteed, lowering incentives for people to move. Often, public providers of employment services (PES) are not well connected either across different locations, preventing jobseekers from getting to know about interesting opportunities. Information sharing is hampered by the lack of a digital infrastructure, outdated modes of information gathering and storage, or simply the use of incompatible standards across different branches of the social protection system. Besides a general investment in the

digital infrastructure, AI-based matching tools could provide instruments to address some of these issues, provided that regulatory barriers to PES are lowered and incentives strengthened to make use of these services. Well-designed social protection systems, therefore, need to include elements of a strong and well-maintained digital infrastructure, portability of rights across occupational and geographical boundaries, and a proper incentive and support structure to help workers in successfully undertaking their transition to a new job opportunity.

Importantly, these social protection systems need to be well funded to provide a sufficient economic stimulus. Indeed, among the difficulties for a successful transition is the lack of a dynamic economic environment to stimulate structural change. Historically, social protection systems have played a significant role in providing insurance against large shortfalls in demand in such situations, and smoothing of aggregate demand thanks to social protection is typically the largest component among labor market policies in contributing to job creation (Ernst, 2015). In this regard, social protection systems can only function when they are supported by well-funded, stable government revenues. Given the increasing importance of superstar firms in the economy, taxing, and redistributing excess profits these firms earn will become increasingly important in ensuring that AI will not lead to an unequal society. Indeed, besides securing funding to social protection systems, an efficient tax system is also an important tool in addressing rising inequality. However, in this age of fast technological change and digitalization, using tax policies to address income inequality faces particular challenges. Technological changes are altering parts of the tax system in important and sometimes dramatic ways, providing both new risks for policymakers and tax administrations to ensure adequate and equal taxation. For instance, digitalization has accelerated the spread of global supply chains in which multinational enterprises integrate their worldwide operations. In this context, taxing rights on income generated from cross-border activities is a challenging task for policymakers. These changes in the age of digitalization and globalization can exacerbate base erosion and profit shifting risks (see OECD, 2015b).

Possible solution to such base erosion is to move from a resident taxation to a customer-based tax system (Falcão, 2018a). Such a system would allow to levy tax revenues where they are generated (i.e., at the level of the individual customer), especially when a large part of the customer base is outside the resident country of the (content) provider. Moving toward a consumption-based tax system is, however, not without its own risks as such taxes might exacerbate income inequalities. In this regard, some recent studies argue that despite their apparent regressive nature, general indirect taxes (e.g., value-added tax, sales tax) can potentially *reduce* income inequality, provided they lead to an increase in labor force participation rates (OECD, 2018; Ciminelli et al., 2017). Nevertheless, indirect taxation of digital contents might need to be complemented with new forms of corporate taxation, which can, when properly designed, stimulate innovation rather than deterring it. As pointed out by Acemoglu et al. (2018), taxing incumbents rather than subsidizing their R&D activities can help strengthen innovation as it will force companies to either innovate or exit the market. Governments can stimulate market exit of low innovative companies, thereby lifting innovation and productivity growth while still ensuring sufficient government revenues. In other words, responding to the needs of properly taxing the digital economy provides opportunities for changes in the tax system that can help maintain a stable tax revenue base while strengthening economy efficiency through higher labor force participation and stronger innovation incentives.

Among the most prominent implications of technological change is that it affects the prices of factors of production (including wages) and of produced goods. A possible way to address both rising income inequalities and skill-biased technological change, therefore, consists in introducing differential taxation to favor labor over capital. Low-skilled workers might benefit, for instance, from wage and hiring subsidies or tax credits, to keep labor demand high for this type of work. Alternatively, tax policies might focus on making capital more expensive, such as the much-discussed robot tax famously advocated by Bill Gates. Such a tax might help to generate significant fiscal revenues without distorting investment incentives, provided that the supply of capital or of inputs complementary to capital are sufficiently inelastic (Korinek and Stiglitz, 2017). More promising solutions include broad resource taxation such as carbon taxes, which would encourage resource-saving instead of labor-saving innovation. It would thus simultaneously address two of the most serious global problems, global climate change and inequality (Falcão, 2018b). Similarly, the elimination of tax deductions for interest and the imposition of a tax on capital would increase the cost of capital and induce more capital augmenting rather than labor-saving innovation. Nevertheless, given the challenges of taxing excess profits arising from digital technologies, alternative ways for a fair distribution of technological rents will need to be considered, which is what we turn to in the next section.

5.4 **How to share technological rents more broadly?**

Rather than trying to tax away excess profits, some policy proposals take issue directly with the way technological rents are currently being appropriated. Indeed, part of the growing inequality produced by the digital economy (and specifically by AI applications) relates to the fact that consumers share their data for free in exchange for "free services." This "zero marginal cost society" was long heralded as the new business model (Rifkin, 2014) but increasingly shows its limitations both in terms of people's privacy concerns and in terms of its economic and social impact, as discussed previously. One solution to address at least the economic side of the issue could be that consumers continue to share their data freely but restrict their use for specific purposes that provide only limited profit opportunities. As soon as a company expects to develop new, profitable products or services – for instance, thanks to medical information that is being shared – consumers' consent needs to be requested and rewarded, for instance, through participation in the expected profits. Given the essential role of data in building and training algorithms for AI tools, such a system could reestablish proper, marginal cost-based incentives in comparison to the current free data–free services model (Ibarra et al., 2018). Such a reward not only rectifies the inequalities that arise from the current system but also maintains incentives for people to share their personal data, a prerequisite for new tools to be designed and developed.

Related to properly setting incentives for data sharing is the issue of data privacy and data control. As mentioned earlier, matching algorithms that rely on large, unstructured databases run the risk of establishing biased profiles of candidates – for instance, on gig or recruitment platforms – that limit employment opportunities and depress working conditions for (certain groups of) candidates, thereby perpetuating preexisting biases. In the case of algorithmic platforms, however, litigation processes are currently underdeveloped or absent, making it difficult to state a case against unfair treatment on these platforms. Several initiatives have already been

started, involving social partners in supporting gig platform workers when faced with such a situation (Berg et al., 2018b). Policymakers and social partners will, however, need to become more alert to these developments, as new applications of AI in areas such as HR analytics are indicating that companies will increasingly cross-analyze large amounts of data in analyzing their workforce performance, some of which is likely to undermine existing national and international labor regulations (De Stefano, 2018). Concrete policy proposals in the use of data for particular purposes will need to be determined and negotiated among social partners on a case-by-case basis, but might involve the restriction of different sources of personal information to be matched for analytical purposes. In this regard, the application and impact of the recently introduced European legislation on General Data Protection Regulation (GDPR) will need to be closely monitored and analyzed to draw useful inferences for other countries and for specific labor market applications.

Several observers have suggested the adoption of policies to share productivity gains more broadly. The two most prominent suggestions are a reduction in working time (possibly combined with a universal basic income support) and shared capital ownership by encouraging workers – either individually or through collective funds – to participate in capital gains and profits ("shared capitalism", Freeman et al., 2009; Freeman, 2016). Neither of these two policy proposals is specific to AI-driven technological change, but given the speed and extent to which AI seems to affect the economy, both proposals can rely on historical experience and might, therefore, easily be implemented and scaled up. A reduction in average working time comes at a moment where productivity gains have not been shared through shorter working weeks over the past few decades but face – in particular in advanced economies – a slowdown or even decrease in labor supply, which might make it difficult to be defended politically. Profit sharing models have also been around as policy proposals for some time and implemented – gradually – among companies and countries in advanced economies (e.g., participation in France). At present, this proposal continues to face strong political resistance, not least because of the fear by capital owners of being restricted in their use of profits and investment. Empirical evidence shows, however, that such a policy could effectively reduce inequalities while at the same time improving company performance (Kurtulus and Kruse, 2017).

Large economies of scale and first-mover advantages from AI (as described above) run the risk of worsening the income gap not only within but also between countries. Convergence achieved over the past three decades by moving people in the developing world out of poverty thanks to increased access to technological transfer and international trade might be put at risk when few companies in advanced economies reap most of the benefits from new, AI-based technologies. In the absence of a fairer international system, many benefits that could accrue to low-income countries thanks to their significantly reduced price of capital might not materialize when leading innovating firms are setting up new barriers to the entry and diffusion of technologies. Many of the potentially development-enhancing AI applications discussed above are developed and patented in advanced economies, leaving access to their use and their benefits among rich nations. Developing countries, therefore, lose out from the benefits of AI on two fronts: first, by not having access to the tax income generated by innovating companies due to the particular way in which international tax treaties allow digital services to be taxed (Falcão, 2018a); and second, as discussed above, by not having an open access to patented AI applications that would be particularly beneficial for their economic development. As patents

are creating a legal monopoly (albeit temporarily), this reinforces the first-mover advantages of digital innovations such as AI, to the detriment of those countries that have less capacity to develop these systems for themselves. Besides open AI approaches previously mentioned, this also calls for action by international development agencies in supporting the implementation of AI and big data strategies in developing countries, helping them to access and develop these technologies for their own national benefit and supporting their diffusion among both private and public actors that would significantly benefit the delivery, implementation and monitoring of policies, such as upholding international labor standards (Grabher-Meyer and Gmyrek, 2017).

6 Outlook and open questions

The current wave of applications based on AI promises to be the largest and most widely ranging technological change observed over the past decades. Its general purpose nature that allows this new technology to be applied in a large span of sectors and occupations, irrespective of the skill level of the involved workforce, creates a broadly shared fear of job loss and control over people's lives. Previous experience with automation, in particular stemming from robotization over the last three decades, seems to suggest that this new wave of technological change might bring significant challenges, especially to developing countries as they face both automation and re-shoring of existing tasks and thereby lose their advantage of lower labor costs that were underpinning their development model over the recent past.

This paper has argued that there are significant opportunities arising from these new, AI-based technologies, including for developing countries, and that the risks, rather than being on the side of job losses, are linked to further worsening income inequalities, both within and across countries. The particular digital nature of AI makes it easy to diffuse but creates large first-mover advantages that can contribute to further rising market concentration and inequality. At the same time, its versatility and general purpose nature allow the creation of expert systems that are potentially beneficial in a large range of occupations, even among low-skilled or low productive ones. In this respect, the paper has argued that the large reduction in capital costs that is brought about by AI applications together with the fact that the direction of technological change is, in part at least, driven by the relative supply of low- vs. high-skilled labor, developing countries stand to benefit from AI, provided it diffuses widely and that technological rents are broadly shared.

For the opportunities to exceed the risks, however, policies need to be adjusted at both the national and the international levels. This paper argues that skills policies in and of themselves, albeit necessary, will not be sufficient in this regard. Policymakers and social partners need to ensure that individual companies cannot gain market dominance, thereby excluding users from their algorithm or maintaining and replicating existing biases. The paper argues that a different way of protecting data is required, giving people more control over their individual information. In addition, existing initiatives such as those undertaken by social partners in the platform economy need to be developed further and implemented more widely. At the international level, a better sharing of the benefits of the new digital economy, possibly through an adjustment in international tax treaties, will also be necessary to prevent digital companies from undermining a country's fiscal revenue base. Finally, long-standing policy proposals for

a fairer global economy should be brought to new life in the light of the significant economic rewards that AI-based innovations promise. This includes a continuous reduction in working hours, especially among those countries where long hours are still the norm, as well as sharing the receipts of innovation rents through profit sharing policies that have already been successfully implemented in some countries in the past.

Given the novelty of this technological innovation, a continuous observation and monitoring of its applications and impact will be necessary, both by national and international actors. Several possible consequences can already be distinguished, such as those discussed in this paper. Others, in particular regarding the specific impact AI-based innovations will have on workplace organization and the employment relationship more broadly, remain highly uncertain. As the technology is evolving quickly, new risks and opportunities might arise that will require constant regulatory adjustment to ensure that technological rents are broadly shared. Also, constant exchanges among policymakers and regulators are necessary to avoid regulatory capture, as well as proper support for local actors to benefit from the advantages of AI. The international community and the ILO, in particular, are well suited to provide this important platform for exchange and experience and to support countries and social partners in adjusting their regulations, as well as negotiation with the necessary information and policy recommendations.

Several questions on the potential long-term consequences of the development of AI arise. One concerns the particular form that AI will take in the future and whether humans will be able to apprehend the decisions being taken by machines. As discussed, current applications of AI run the risk of replicating biases from human decisions (e.g., in hiring). This poses obvious ethical questions, in particular as these applications no longer allow a transparent account of how the decisions have been taken. Recent developments in this area that rely on a different methodology (genetic algorithms rather than neural networks) might offer a more transparent alternative, but for the moment it is too early to assess their full potential. Another, more fundamental, question concerns the shift from automating workforce to automating "brain force," with machines (autonomously) acquiring new skills and competencies at a much faster pace than humans will be able to. If such a shift from specific AI (as discussed in this paper) to general AI takes place, human capital will no longer be the constraining factor in the technological evolution, which could happen much faster than before. In other words, evolution would no longer be constrained by "biological computers" (i.e., humans) but could move to machines, a vision recently put forward by Harari (2016). Ultimately, however, the type of AI algorithms that will be used and the decision to develop and implement general AI – independently of its technical feasibility – will eventually be determined by policymakers and customers who might deliberately vote and decide against some of the more harmful manifestations of AI. For the moment, at least, it remains the case that robots cannot vote.

Declarations

Availability of data and materials

The datasets used and/or analyzed during the current study are available from the corresponding author on reasonable request.

Competing interests

The authors declare that they have no competing interests.

Funding
No external funding was provided or used in preparing this paper.

Authors' contributions
The paper was organized and drafted by EE. RM provided an initial draft of Sections 5.2, 5.3, and 5.4. DS provided initial thoughts and draft inputs into Sections 1 and 3. Both RM and DS reviewed the draft and made comments and corrections to the draft.

Acknowledgments
Research assistance, in particular for Section 2, by Francesco Carbonero (ITC, Turin, Italy) is gratefully acknowledged. Earlier drafts of this paper have been reviewed by Prof. James Bessen (Boston University), Lisa Feist, Pawel Gmyrek, L. Jeff Johnson, Hannah Johnston, Irmgard Nübler (all ILO), Prof. Enzo Weber (IAB, University of Regensburg); their comments and suggestions provided helpful input to prepare this paper. We also thank the journal editor, Prof. Denis Fougère, and an anonymous referee for very helpful comments. All remaining errors are ours.

References
Acemoglu, D. (2002): Technical Change, Inequality, and the Labor Market. Journal of Economic Literature 40(1), 7-72.
———; U. Akcigit; H. Alp; N. Bloom; W. Kerr (2018): Innovation, Reallocation, and Growth. American Economic Review 108(11), 3450-3491.
———; D. Autor (2011): Skills, Tasks and Technologies: Implications for Employment and Earnings, in: Ashenfelter, O.; D. Card (eds.), Handbook of Labor Economics, Vol. 4B. Amsterdam, North Holland: Elsevier, 1043-1172.
———; P. Restrepo (2017): Robots and Jobs: Evidence from US Labor Markets. NBER Working Paper No. 22252. Cambridge, MA: National Bureau of Economic Research.
———; (2018): Artificial Intelligence, Automation and Work. NBER Working Paper No. 24196. Cambridge, MA: National Bureau of Economic Research.
Aghion, P.; N. Bloom; R. Blundell; R. Griffith; P. Howitt (2005): Competition and Innovation: An Inverted-U Relationship. Quarterly Journal of Economics 120(2), 701-728.
Agrawal, A. K.; J. S. Gans; A. Goldfarb (2018a): Prediction Machines: The Simple Economics of Artificial Intelligence. Boston, MA: Harvard Business Review Press.
———; (2018b): Economic Policy for Artificial Intelligence. NBER Working Paper No. 24690. Cambridge, MA: National Bureau of Economic Research.
Albertini, J.; J.-O. Hairault; F. Langot; T. Sopraseuth (2017): A Tale of Two Countries: A Story of the French and US Polarization. IZA Discussion Paper No. 11013. Bonn: Institute for the Study of Labor.
Andrews, D.; C. Criscuolo; P. N. Gal (2016): The Best Versus the Rest: The Global Productivity Slowdown, Divergence Across Firms and the Role of Public Policy. OECD Productivity Working Paper No. 5. Paris: Organisation for Economic Co-operation and Development.
Arntz, M.; T. Gregory; U. Zierahn (2016): The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis. OECD Social, Employment and Migration Working Paper No. 189. Paris: Organisation for Economic Co-operation and Development.
———; (2017): Revisiting the Risk of Automation. Economic Letters 159, 157-160.
Autor, D. H. (2010): The Polarization of Job Opportunities in the U.S. Labor Market: Implications for Employment and Earnings. Washington, DC: Center for Economic Progress.
———; D. Dorn; L. Katz; C. Patterson; J. Van Reenen (2017a): Concentrating on the Decline in Labor's Share. American Economic Review Papers and Proceedings 107(5), 180-185.
———; (2017b): The Fall of the Labor Share and the Rise of Superstar Firms. NBER Working Paper No. 23396. Cambridge, MA: National Bureau of Economic Research.
———; F. Levy; R. J. Murnane (2003): The Skill Content of Recent Technological Change: An Empirical Exploration. Quarterly Journal of Economics 118(4), 1279-1333.
———; L. F. Katz; M. S. Kearney (2006): The Polarization of the U.S. Labor Market. American Economic Review Papers and Proceedings 96(2), 189-194.
Bárány, Z. L.; C. Siegel (2018): Job Polarization and Structural Change. American Economic Journal: Macroeconomics 10(1), 57-89.
Benhamou, S.; L. Janin (2018): Intélligence artificielle et travail. Paris: France Stratégie.
Berg, A.; E. F. Buffie; L.-F. Zanna (2018a): Robots, Growth, and Inequality: Should We Fear the Robot Revolution? (The Correct Answer is Yes). IMF Working Paper No. 18/116. Washington DC: International Monetary Fund.

Berg, J.; M. Furrer; E. Harmon; U. Rani; M. S. Silberman (2018b): Digital Labour Platforms and the Future of Work: Towards Decent Work in the Online World. Geneva: International Labour Office.

Bessen, J. (2015a): How Computer Automation Affects Occupations: Technology, Jobs, and Skills. Law and Economics Research Paper No. 15-49. Boston: Boston University School of Law.

———; (2015b): Learning by Doing: The Real Connection Between Innovation, Wages, and Wealth. New Haven: Yale University Press.

———; (2017a): Information Technology and Industry Concentration. Law and Economics Research Paper No. 17-41. Boston: Boston University School of Law.

———; (2017b): Automation and Jobs: When Technology Boosts Employment. Law and Economics Research Paper No. 17-09. Boston: Boston University School of Law.

———; (2018): AI and Jobs: The Role of Demand, Mimeo. http://www.nber.org/chapters/c14029.pdf. Accessed 13 June 2018.

Boone, J. (2001): Intensity of Competition and the Incentive to Innovate. International Journal of Industrial Organization 19, 705-726.

Braverman, H. (1974): Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century. New York: Monthly Review Press.

Bunker, N. (2016): Why Declining US Labour Mobility is About More Than Geography. New York: World Economic Forum. https://www.weforum.org/agenda/2016/09/why-declining-us-labour-mobility-is-about-more-than-geography.

Carbonero, F.; E. Ernst; E. Weber (2018): Robots Worldwide: The Impact of Automation on Employment and Trade, ILO Research Department Working Paper No. 36. Geneva: International Labour Office.

Chiacchio, F.; G. Petropoulos; D. Pichler (2018): The Impact of Industrial Robots on EU Employment and Wages: A Local Labour Market Approach. Bruegel Working Paper No. 2. Brussels, Bruegel. http://bruegel.org/wp-content/uploads/2018/04/Working-Paper-AB_25042018.pdf. Accessed 14 June 2018.

Ciminelli, G.; E. Ernst; M. Giuliodori; R. Merola (2017): The Composition Effects of Tax-Based Consolidations on Income Inequality. ILO Research Department Working Paper No. 19. Geneva: ILO.

Cockburn, I. M.; R. Henderson; S. Stern (2018): The Impact of Artificial Intelligence on Innovation. NBER Working Paper No. 24449. Cambridge, MA: National Bureau of Economic Research.

Cohen, R. B. (2018): Business Use of Artificial Intelligence (AI) and Machine Learning (CL) and its Impact on the U.S. Economy. ESI Research Report. Washington DC: Economic Strategy Institute.

Culey, S. (2012): Transformers: Supply Chain 3.0 and how Automation will Transform the Rules of the Global Supply Chain. The European Business Review (Online).

Danninger, S. (2016): What's up with U.S. Wage Growth and Job Mobility? IMF Working Paper No. 16/122. Washington DC: International Monetary Fund.

David, P. A. (1990): The Dynamo and the Computer: A Historical Perspective on the Modern Productivity Paradox. American Economic Review Papers and Proceedings, 355-361.

De Backer, K.; T. DeStefano; C. Menon; J. R. Suh (2018): Industrial Robotics and the Global Organisation of Production. OECD Science, Technology and Industry Working Paper No. 2018/03. Paris: Organisation for Economic Co-operation and Development.

———; **C. Menon; I. Desnoyers-James; L. Moussiegt** (2016): Reshoring: Myth or Reality? OECD Science. Technology and Industry Working Paper No. 2016/27. Paris: Organisation for Economic Co-operation and Development.

Deming, D. J. (2017): The Growing Importance of Social Skills in the Labor Market. Quarterly Journal of Economics 132(4), 1593-1640.

De Stefano, V. (2018): "Negotiating the Algorithm": Automation, Artificial Intelligence and Labour Protection. Employment Working Paper No. 246. Geneva: International Labour Office.

Ernst, E. (2015): How Unemployment Benefits can Help Unemployed Workers and Strengthen Job Creation. International Social Security Review 68(3), 43-67.

———; (2018): Is Technological Change Accelerating? ILO Briefing for the Global Commission on the Future of Work. Geneva: International Labour Office.

———; **Chentouf, L.** (2014): Work Organisation and Incentives. Global and Local Economies Review 18(1), 103-135.

Falcão, T. (2018a): Taxing the Digital Economy: Policy Considerations and how to Advance the Debate. Tax Notes International 12 Feb, 623-628.

———; (2018b): "Should My Dishwasher Pay A Robot Tax? Tax Notes International 11 June, 1273-1277.

Frey, C. B.; M. A. Osborne (2017): The Future of Employment: How Susceptible are Jobs to Computerisation? Technological Forecasting and Social Change 114(C), 254-280.

Freeman, R. (2016): Who Owns the Robots Rules the World: The Deeper Threat of Robotization. Harvard Magazine May-June. https://harvardmagazine.com/2016/05/who-owns-the-robots-rules-the-world.
———; **J. Blasi; D. Kruse** (2009): Shared Capitalism at Work: Employee Ownership, Profit and Gain Sharing, and Broad-Based Stock Options. Chicago: University of Chicago Press.
Fujii, H.; S. Managi (2018): Trends and Priority Shifts in Artificial Intelligence Technology Invention: A Global Patent Analysis. Economic Analysis and Policy 58(C) 60-69.
Furman, J.; R. Seamans (2018): AI and the Economy. NBER Working Paper No. 24689. Cambridge, MA: National Bureau of Economic Research.
Gifford, D. J.; R. T. Kudrle (2010): "The Law and Economics of Price Discrimination in Modern Economies: Time for Reconciliation? UC Davis Law Review, No. 1235. http://scholarship.law.umn.edu/faculty_articles/358.
Goldin, C.; L. Katz (1998): The Origins of Technology-Skill Complementarity. Quarterly Journal of Economics 113(3), 693-732.
Gordon, R. J. (2016): The Rise and Fall of American Growth. The U.S. Standard of Living Since the Civil War. Princeton, NJ: Princeton University Press.
Grabher-Meyer, N.; P. Gmyrek (2017): Big Data and Artificial Intelligence in the ILO's Development Cooperation. State of Play, Ideas, Opportunities and Threats. Geneva: ILO.
Graetz, G.; G. Michaels (2015): Robots at Work. IZA Discussion Paper No. 8938. Bonn: Institute for the Study of Labor.
Harari, Y. N. (2016): Homo deus. A Brief History of Tomorrow. London: Harvill Secker.
Ibarra, I. A.; L. Goff; D. J. Hernández; J. Lanier; E. G. Weyl (2018): Should we Treat Data as Labor? Moving Beyond 'Free'. American Economic Association Papers and Proceedings 108, 38-42.
ICF. (2016): Global Coaching Study. Lexington, KS: International Coach Federation.
ILO. (2015): World Employment and Social Outlook. Trends. Geneva: International Labour Office.
———; (2016): Global Wage Report: Wage Inequality in the Workplace. Geneva: International Labour Office.
———; (2018): World Employment and Social Outlook. Trends. Geneva: International Labour Office.
IMF. (2018): Technology and the Future of Work. G20 Background Note. Washington DC: International Monetary Fund. http://www.imf.org/external/np/g20/pdf/2018/041118.pdf. Accessed 13 June 2018.
Kearney, A. T. (2017): Global Services Location Index: The Widening Impact of Automation. New York. https://www.atkearney.com/documents/20152/793366/The+Widening+Impact+of+Automation.pdf.
Korinek, A.; D. X. Ng (2017): The Macroeconomics of Superstars. Mimeo. http://www.korinek.com/download/Superstars.pdf.
———; **J. Stiglitz** (2017): Artificial Intelligence and its Implications for Income Distribution and Unemployment. NBER Working Paper No. 24174. Cambridge, MA: National Bureau of Economic Research.
Kurtulus, F.; D. Kruse (2017): How did Employee Ownership Firms Weather the Last Two Recessions? Employee Ownership, Employment Stability, and Firm Survival: 1999-2011. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.
Logg, J. M.; J. A. Minson; D. A. Moore (2018): Algorithm Appreciation: People Prefer Algorithmic to Human Judgment. Harvard Business School Working Paper No. 17-086. Cambridge, MA: Harvard University.
Marglin, S. A. (1974): What do Bosses do? The Origins and Functions of Hierarchy in Capitalist Production. Review of Radical Political Economics 6(2), 60-112.
Marshall, A. (1890): Principles of Economics: London: Macmillan.
Mazzucato, M. (2013): The Entrepreneurial State: Debunking Public vs. Private Sector Myths. London: Anthem Press.
Méda, D. (2016): The Future of Work: The Meaning and Value of Work in Europe. ILO Research Paper No. 18. Geneva: International Labour Office.
MGI (2015): A Labor Market that Works: Connecting Talent with Opportunity in the Digital Age. Washington DC: McKinsey Global Institute.
———. (2017): Reinventing Construction: A Route to Higher Productivity. Washington DC: McKinsey Global Institute.
———. (2018a): AI, Automation, and the Future of Work: Ten Things to Solve for. Washington DC: McKinsey Global Institute.
———; (2018b): Skill Shift. Automation and the Future of the Workforce. Washington DC: McKinsey Global Institute.
Molloy, R.; C. S. Smith; A. K. Wozniak (2014): Declining Migration within the US: The Role of the Labor Market. NBER Working Paper No. 20065. Cambridge, MA: National Bureau of Economic Research.
Moretti, E. (2012): The New Geography of Jobs. Boston: Houghton Mifflin Harcourt.
Naudé, W.; P. Nagler (2015): Industrialisation, Innovation, Inclusion. UNIDO Working Paper No. 15/2015. Vienna: United Nations Industrial Development Organization.

OECD (2015a): Data-Driven Innovation. Paris: Organisation for Economic Co-operation and Development.
———. (2015b): Addressing the Tax Challenges of the Digital Economy ACTION 1: 2015 Final Report. Paris: Organisation for Economic Co-operation and Development.
———; (2017): Digital Economy Outlook. Paris: Organisation for Economic Co-operation and Development.
———. (2018): Tax Challenges Arising from Digitalisation: Interim Report 2018: Inclusive Framework on BEPS, in: OECD/G20 Base Erosion and Profit Shifting Project. Paris: Organisation for Economic Co-operation and Development.
Ponce Del Castillo, A. (2018): Artificial Intelligence: A Game Changer for the World of Work. Foresight Brief No. 5. Brussels: European Trade Union Institute.
Rifkin, J. (2014): The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism. New York: St. Martin's Press.
Rosen, S. (1981): The Economics of Superstars. American Economic Review 71(5), 845-858.
Runst, P. (2018): The Effect of Occupational Licensing Deregulation on Migrants in the German Skilled Crafts Sector. European Journal of Law and Economics 45(3), 555-589.
Scassa, T. (2018): Data Ownership. CIGI Papers No. 187. Waterloo: Centre for International Governance Innovation.
Schwab, K. (2016): The Fourth Industrial Revolution. Geneva: World Economic Forum.
Semmler, W.; P. Chen (2017): Short and Long-Run Effects of Productivity on Unemployment, Mimeo. https://ssrn.com/abstract=2907539. Accessed 12 June 2018.
Sengenberger, W. (1987): Struktur und Funktionsweise von Arbeitsmärkten: Die Bundesrepublik Deutschland im internationalen Vergleich. Frankfurt am Main: Campus Verlag.
Simon, H. (1960): The New Science of Management Decision. New York: Harper.
Sirkin, H. L.; M. Zinser; J. R. Rose (2014): The Shifting Economics of Global Manufacturing: How Cost Competitiveness is Changing Worldwide. Boston, MA; Boston Consulting Group.
Stucke, M. E.; A. P. Grunes (2016): Big Data and Competition Policy. Oxford: Oxford University Press.
Tirole, J. (1988): The Theory of Industrial Organization. Cambridge, MA: MIT Press.
Turing, A. M. (1950): Computing Machinery and Intelligence. Mind 49, 433-460. https://www.csee.umbc.edu/courses/471/papers/turing.pdf.
UNCTAD. (2016): Robots and Industrialization in Developing Countries. Policy Brief No. 50. Geneva: United Nations Conference on Trade and Development.
Varian, H. R. (1985): Price Discrimination and Social Welfare. American Economic Review 75(4), 870-875.
Vivarelli, M. (2014): Innovation, Employment and Skills in Advanced and Developing Countries: A Survey of Economic Literature. Journal of Economic Issues 48(1), 123-154.`World Bank (2016): World Development Report: Digital Dividends. Washington, DC.
Zhang, X.; J. Yang; S. Wang (2011): "China has Reached the Lewis Turning Point. China Economic Review 22(4), 542-554.