econfip
Economists for Inclusive Prosperity

Labor in the Age of Automation and Artificial Intelligence

Anton Korinek*

As technology advanced in recent decades, it increasingly left workers behind and led to sharp increases in inequality. The current wave of progress in artificial intelligence is likely to accelerate these trends. This note lays out three complementary approaches to countering these developments. Firstly, since technological progress generates net gains for society as a whole, the winners could in principle compensate the losers and still be better off. Secondly, progress should be steered to minimize the losses of workers. Thirdly, there is an important role for government intervention in information technology to thwart the rise of monopolies that extract rents from society. The note concludes with some speculations on the impact of artificial intelligence increasingly rivaling human labor.

Introduction

By and large, workers have had a good run over the first two centuries of the Industrial Revolution. Technological progress automated tasks involving hard physical labor and tedious routine activities and, as a by-product, increased workers' incomes about ten-fold. In the eyes of idealists, progress freed human workers from tasks that were in fact inhumane, that humans were never meant to perform – so they can focus on more fulfilling work involving cognitively interesting activities.

Starting about four decades ago, however, technological progress has increasingly left workers behind, as reflected in a range of dismal statistics: Since World War II, the labor force participation rate of prime-age men has declined from 98% to 89%; and the labor share of income – what is earned by workers rather than capitalists – has declined from 66% to 58% (see left panel of the figure 1). The average real wage of regular workers has in fact declined over the past four decades – over a period in which total income in the US almost tripled! At the same time as regular workers fell behind, the so-called superstars of the economy have garnered an increasing share of income, with the top 1% more than doubling their take to about 20% of all income, and the top 0.1% tripling their take to close to 10% (see right panel of the figure 1). Looking at wealth rather than income, some estimates suggest that the richest three Americans now own more than what the bottom 50% of the US population own.

Although technological forces were undeniably a prime force behind these developments, there were also other factors involved, many of which are discussed

Figure 1

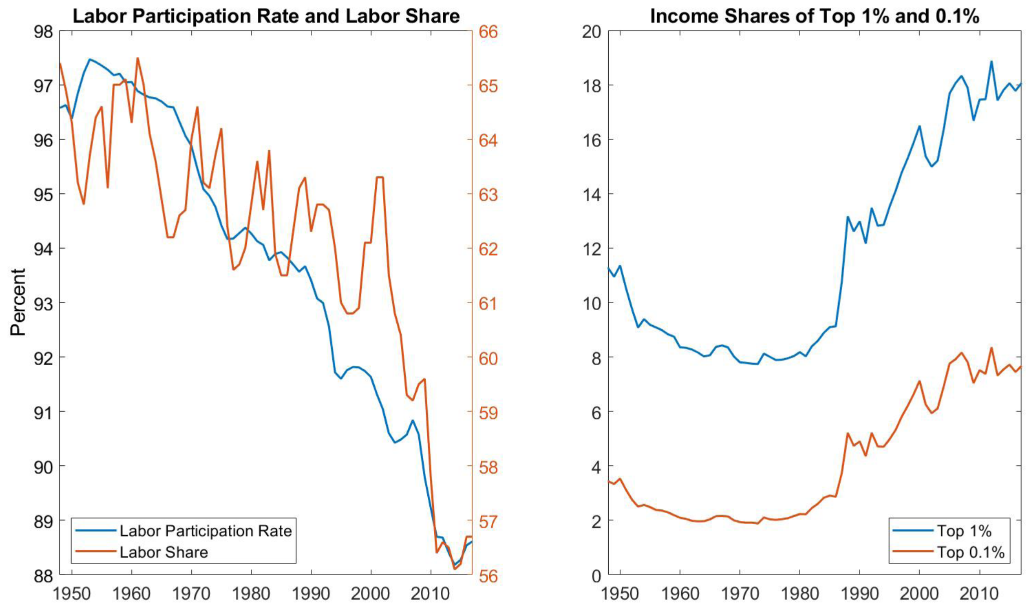

comprehensively by other policy briefs of this series by Economists for Inclusive Prosperity. For example, trade liberalization put pressure on workers competing with cheaper labor abroad. Institutional changes such as less generous redistributive policies, the declining power of unions, and tax policies favoring the rich reduced the take-home income of regular workers. Many of these factors were in fact also partly driven by technological forces.

This brief focuses squarely on the implications of technological change and how to manage them. The traditional approach of economists has been to view technology as a driving force that is outside the realm of our analysis – technology is developed by engineers; we economists take it as given and study the implications for the economy. But technology is not destiny. In fact, better understanding the technological forces behind the decline of labor is crucial for shaping our agenda on how to best protect workers going forward. In the following, I will discuss the broader forces that have contributed to rising inequality in recent decades and how to counter-act them; I will zoom in on the implications of information goods and digital technologies for the economy; finally, I will speculate on how the rise of artificial intelligence will affect workers and the economy in coming decades. In each section, I will include a discussion of the policy options available.

Does Technological Progress Replace Workers?

It has been true since the advent of the Industrial Revolution that technological progress replaced specific jobs – at the time, for example, spinners and weavers. But time and time again, after a period of adjustment, the economy created new jobs for the displaced workers that ultimately paid better. (If the disruption was severe, the adjustment took longer, but at least the children of displaced workers found better jobs.) Many economists therefore proclaim that technological progress is unambiguously good for workers.

However, it is not a natural law that technological progress will lead to higher wages and improved livelihoods. The fact that real wages of regular workers have declined over the past four decades, strongly suggests that the overall effect of technological *change* over that period (one hesitates to call it technological progress) has been to reduce the market wages of regular (unskilled) workers. As technology advanced in recent decades, the economy simply seemed to have less and less need for unskilled labor.

At first sight, the picture looks better for skilled workers who saw their wages rise significantly over the 1980s and

Figure 2

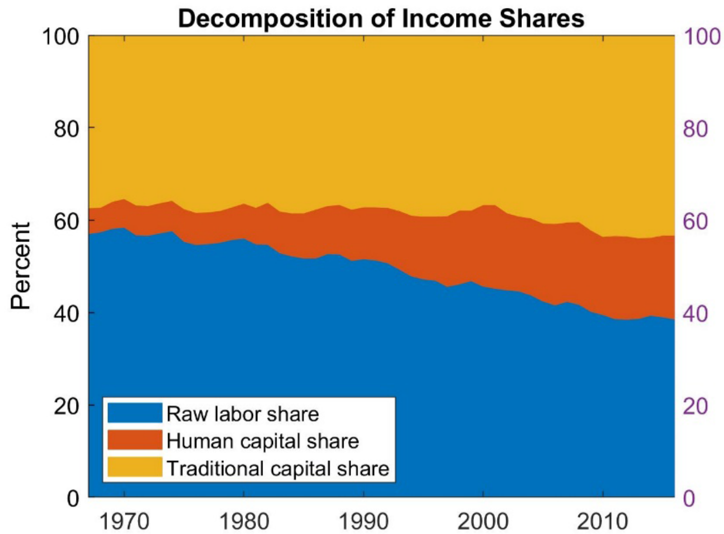

1990s, although less so in the recent two decades (see e.g. Autor, 2015; Brinca et al., 2019). However, there is also a bleaker interpretation of this phenomenon: skilled labor can be interpreted as unskilled labor enhanced by education, i.e. it is a composite of unskilled labor and human capital. The wages of skilled workers can, in this view, be decomposed into the wages of unskilled workers plus returns on human capital. Since the wages of the unskilled have not increased, all of the increase in skilled wages in fact reflects returns on human capital investment.

Figure 2 above decomposes what fraction of national income was earned by unskilled labor versus capital (made up of both traditional and human capital) from 1967 to 2017. The "human capital share," calculated as the extra returns earned by college graduates, has risen from 5.6% to 18.2% of total income. Conversely, the raw labor share has declined from 57% to less than 40%. According to this interpretation, combined capital earns more than three fifths of all output in the economy.

From this point of view, the past four decades have led to an even starker reallocation of returns from labor to capital (which now includes human capital). The difference matters for workers, since wage earnings reflect the return on human labor effort, whereas the returns on human capital are returns on investments in education, which is becoming ever more costly and which in fact soaks up a large part of these returns.

Technological Redistribution and Social Redistribution

When technological progresses leads to income losses for workers, it is natural to think about ways to compensate them for their losses. An important point to emphasize is that technological progress could in principle make everybody better off, i.e. generate what economists call a Pareto improvement, if there is sufficient political will. By definition, technological progress means that the economy can produce more for a given amount of inputs, implying that there is more overall income to be distributed. If one of the factors of production, for example unskilled labor, earns less as a result of an innovation, it means that someone else is not only earning the additional fruits of progress, but also appropriating part of what used to be the earnings of unskilled labor.

More generally, we can decompose the economic effects of technological progress into two parts, as laid out e.g. in Korinek and Stiglitz (2019): First, progress raises overall output, i.e. it increases the size of the economic pie produced. This extra output is earned by someone in the economy, for example by the innovator who reaps the fruits of her innovation. Secondly, technological progress also generates a redistribution of the existing economic pie, as it changes the market prices at which people transact in the economy. For example, it may reduce the wages of some workers and increase the wages of other workers. This redistribution via price changes is always zero-sum: price increases benefit sellers at the expense of buyers, and vice versa for price reductions. We may call this effect a technological redistribution of income since it is generated by technological forces playing out in the market economy.

Consider for example a new AI system that replaces human radiologists: the first effect of such a system may be to lead to better diagnosis and increased use of radiology services, increasing the size of the economic pie produced. The second effect may be to reduce the market wages of human radiologists who had specialized in interpreting images but to increase the wages of nurses and of specialists who rely on radiologists, who can now do more without requiring input from costly human radiologists.

In an idealized world, if we want to avoid that technological progress leaves behind some members of society, then social redistribution would aim to undo the described technological redistribution to compensate the losers of innovation while keeping the increased size of the pie. Those who benefit from technological redistribution accrue windfall gains, i.e. gains that are not based on their own efforts but more on luck. From a policy perspective, this makes it important to be explicit about who are the beneficiaries of technological redistributions, and to look for ways to tax their windfall gains to compensate the losers. At times, it may be possible to tax away such windfall gains without introducing the distortions that taxes usually generate. In those cases, undoing technological redistribution may be feasible without any efficiency losses to the economy. For example, if an innovation increases the value of land in particular areas, higher property taxes could tax away these windfall gains. Korinek and Stiglitz (2018) show that this is a more efficient solution than e.g. the "robot taxes" that have been proposed.

However, there are also many technological redistributions that are quite difficult to undo in practice. In our radiology example above, we would need to tax away the wage gains of nurses and other specialists to compensate radiologists for their losses. This is a proposition that is impractical because of a variety of information problems – it would require a system of taxes and transfers that is far more fine-tuned than what is possible, and it would introduce a number of moral hazard problems, for example reducing work incentives for the beneficiaries of technological redistribution who would face higher tax burdens. Furthermore, some may also view it as unfair if nurses are taxed to compensate better-earning radiologists for the losses stemming from technological redistribution. In all those cases, a general progressive tax system – one that charges higher rates to individuals earning more – together with a social safety net that limits the downside for the losers of technological progress may be one of the best available second-best solutions.

Steering technological progress

Technological progress is the result of conscious and targeted efforts of innovators – unlike the way it is described in many of our economic models, in which progress just happens exogenously. When an innovator comes up with ingenious new methods of producing novel goods or services, or with novel processes to produce existing goods or services more cheaply, her incentives are set by market forces – but the price signals sent by the market do not always reflect social value. This phenomenon is well-known when it comes to externalities such as pollution, and there is wide consensus among economists to regulate such externalities and correct the price signals sent by the market to better reflect our social values.

If we as a society care about inequality and we cannot realistically achieve the desired income distribution purely via transfers, then it is natural to extend this framework of correcting price signals to make innovators more conscious of the distributive externalities of their inventions. For example, if an innovator comes up with a clever new technology to replace the work of thousands of unskilled workers with a handful of skilled workers, the innovation will create a large technological redistribution – unskilled worker will see their wages decline and skilled workers will experience wage gains, exacerbating the trends of the past four decades. Neither of the two groups of workers have actively contributed to these windfall gains and losses, so they constitute externalities. Economists have traditionally been skeptical of such arguments because the described externalities are so-called "pecuniary" externalities – they occur because market prices and wages adjust. If we only care about efficiency not equity, then it is desirable to ignore pecuniary externalities. If we also care about equity, then pecuniary externalities are at the center stage and need to be addressed to achieve our goals.

Technology policy should thus steer technological progress so as to encourage innovation that has desirable distributive properties and to discourage innovations that increase inequality. Let me outline three different avenues for doing this:

The first avenue is to focus on the distributive implications of all the research that is conducted or sponsored by government. Government is one of the largest sources of research funds in our economy, and it should actively steer progress in directions that augment workers rather than replacing them. One example of this is what has come to be called *intelligence assistance*, i.e. AI systems that are designed to complement and enhance the abilities of workers so they can perform higher value-added tasks. Such intelligence assistance may make it possible for workers to do jobs that were previously out of reach for them, greatly increasing the demand for unskilled labor. If intelligence assistance systems were privately funded, there is significant risk that their creators will reap most of the economic returns; if they are publicly funded, by contrast, they can be made available for free or at cost, and workers can reap the resulting returns.

The second avenue is to use regulatory powers as well as tax and subsidy schemes to steer technological progress, in a similar fashion to how other types of externalities such as pollution are addressed. If it is possible to identify whether a specific type of innovation will have positive or negative distributive effects, then the innovative activity itself could be subsidized or taxed, or patent lives on the respective innovations could be lengthened or shortened. Otherwise, subsidizing the employment of lower-skilled workers would lower the cost of such employment and provide socially more desirable price signals to innovators (just like putting a price on carbon induces innovation to engage in carbon-saving activities). For example, if unskilled labor becomes cheaper, then it is less desirable to develop innovations that save on unskilled labor.

A third avenue to steering the path of technological progress is simply to create more awareness of the distributive implications of different types of innovation. Although it is difficult to predict what the exact impact of an innovation on labor markets will be, there are some general guidelines: for example, process innovations that reduce costs by automating labor are more likely to hurt workers than product innovations that generate new products that meet previously unknown needs. Many entrepreneurs are socially-minded and care about the impact of their innovations on society. Making them more aware of the distributive implications of their actions will make a difference. There is also a vibrant NGO sector in the US, partly funded by high tech billionaires, that could make it one of its priorities to invest in innovations such as intelligence assistance that complement unskilled workers rather than displacing them.

Digitization, information goods, and the rise of superstars

An aspect of the recent wave of technological progress that sets it apart from earlier waves of progress is that it centers on digitization and information goods. This is most visible in the IT sector, where some companies generate billions of dollars of revenue selling digital goods while employing just a few hundred employees (who are usually highly skilled) to produce them. However, information goods are a broad phenomenon that is increasingly relevant not only in the technology sector but throughout the economy: in sectors from retail to the food and beverage industry, productive companies such as Walmart or Starbucks replicate their success in local market after market by copying an information good – best business practices – over and over again.

What makes information different from tangible goods is that it is *non-rival* but *excludable*. Non-rival means that it can be used without being used up: if somebody writes a computer program, billions of people can use the same code without using it up. By contrast, tangible goods are typically rival and are eventually used up when they are used: if someone eats a loaf of bread, no one else can eat it. The non-rival nature of information goods implies that once a company has incurred the cost of developing them, it can copy them many times at negligible marginal cost. This means that sectors in which information goods play an important role are *natural monopolies*: it is most economical to develop an information good only once or (if tastes differ) a small number of times, and then to distribute it to the entire market.

If an information good is created by a private owner, the excludable nature of such goods implies that its owner can prevent competitors from using it and therefore has market power. This enables the owner to charge higher prices and extract monopoly rents. In our paper on "The Macroeconomics of Superstars" (with Ding Xuan Ng, 2018), we argue that most of the rise of market power in recent decades and the associated rents can be explained by digitization and information goods across the US economy. This has also been an important factor behind the rise in inequality over the period.

Public policy faces two fundamental problems when dealing with information goods:

The first fundamental problem is that the private market has difficulty achieving efficient outcomes when information goods are involved. One the one hand, financing information goods necessitates that private companies have some monopoly power so they can charge a markup over their marginal cost and earn rents to recoup the cost of their investment. Our society typically grants such monopoly power by awarding intellectual property rights to the creators of information goods that provide them with exclusivity. On the other hand, the monopoly markups that such firms are charging imply that consumers face higher prices and will demand less than what is efficient. As a result, the private market will both underprovide and underuse information goods. Furthermore, it turns those who successfully commercialize information goods into so-called superstar firms, leading to large increases in inequality.

The most efficient solution in the face of these problems would be to publicly fund the creation of information goods and then – since they are almost free to copy – distribute them at a very low price (technically, at marginal cost) to anybody who is interested in using them. This model works relatively well for fundamental research. An example is when DARPA used public funds to finance the invention of the Internet, which has since created trillions of dollars of value. The role of government in financing information goods and making them freely available to society should be expanded as much as possible. Making information goods available for free also has positive distributive implications as it avoids the large monopoly rents that otherwise accrue to the holders of intellectual property rights.

However, the second fundamental problem is that when it comes to commercializing products, private companies are frequently superior to publicly funded entities. For example, Steve Jobs was probably better at designing iPhones than DARPA would have been. Even if DARPA wanted to, it could not contract out the design of goods that have not yet been imagined to visionaries such as Steve Jobs. This limits the spheres in which the efficient outcome can be achieved via public investment in information goods.

In all areas where we rely on private actors for commercialization, we are thus left with second-best policy options that involve granting some monopoly power to the private creators of information goods by awarding them intellectual property rights. When dealing with second-best policy options, everything is about trade-offs: Although granting limited monopoly power may be desirable to provide incentives to innovate, the level of monopoly power that we currently provide and the resulting monopoly rents seem far in excess of the cost of investment in a number of industries, as indicated by record profits. Since the resulting rents extract surplus from consumers to the benefit of large corporations, with undesirable distributive implications, we should counteract them. One avenue is to weaken intellectual property rights and the associated monopoly power; another avenue is to tax away some of the rents earned by corporations, e.g. by charging them licensing fees for the publicly created technologies that they rely on.

Increasing returns that stem from network externalities are an additional factor that is very relevant in the context of digitization and information goods. The greater the number of existing users on a digital platform such as Facebook, Google or Amazon, the more attractive the platform becomes for new users. This makes platforms even stronger natural monopolies, amplifying the associated rents and superstar effects as well as the resulting inequality. Since these network effects frequently revolve around data, we can reduce the power of such natural monopolies by giving consumers more freedom in how their data is used and by forcing inter-operability between different platforms via standards for data exchange. For example, if consumers can grant a start-up that they trust access to the same social network graph, search history or shopping history that established internet firms already have, the monopoly power of existing corporations would be curbed.

The Rise of Artificial Intelligence

Digitization and superstar firms are just the beginning of a larger wave of technological progress that will be of increasing relevance going forward and that centers on the rise of Artificial Intelligence (see e.g. Agrawal et al., 2019; Korinek and Stiglitz, 2019; Acemoglu and Restrepo, 2019). Traditionally, when we were concerned about inequality, we have been thinking of inequality between different groups of humans, such as workers and capitalists, and how they compete for scarce resources. This is based on the anthropocentric notion that only humans consume final goods – a notion that has been perfectly reasonable for much of the history of mankind.

At the present stage, we humans still feel mostly in control of the intelligent algorithms and machines that we have created and that we interact with on a daily basis. However, to an objective observer, things look a little bit different than they used to a few decades ago: Artificially intelligent agents (AIAs) play an increasingly important role in our economy and are in fact more and more in control of us humans. A growing number of corporate decisions that affect us are made by AIAs – from screening job applicants to providing loans. A growing number of our personal decisions is strongly influenced (or, one might say, manipulated) by AIAs – ranging from what we read and buy to whom we date and how we vote. AIAs also act increasingly autonomously in our economy, for example engaging in financial transactions or driving on our roads.

From a broader perspective, humans and intelligent machines are both entities that share certain basic economic properties: first and foremost, they both absorb scarce resources. These resources serve to meet their maintenance needs and ensure their survival. Although the absorption basket of the two types differ – for example, humans consume bread whereas machines absorb electricity – the basic economic function is the same. Furthermore, both types of entities also supply their factor services to the economy – human labor or machine labor, and they both follow defined laws-of-motion.

In a recent paper on "The Rise of Artificially Intelligent Agents" (Korinek, 2018), I observe that – as we are entering a period in which ever more intelligent machines surpass the capabilities of humans in a growing number of areas – competition over scarce resources may increasingly play out between humans and artificial entities. The most tangible present manifestation of such entities are high-tech corporations. They absorb a growing share of the economy's resources – for example, the human labor they employ, the raw materials that go into computing and data centers, and the electricity they consume (server farms absorb close to 10% of the world's electricity production, by some estimates). They also accumulate rising levels of wealth. And although they are notionally owned by humans, the de-facto level of control exerted by their owners is rather low. As AIAs gain ever more autonomy, their actions increasingly surprise their human creators and owners and are frequently misaligned with the objectives of their human owners, as for example Mark Zuckerberg experienced when he found out about the role of Facebook in recent elections. From this perspective, the question of "who owns intelligent machines or algorithms" is increasingly irrelevant – ownership without control is meaningless. The true masters of the universe, as Silicon Valley refers to the founders of the largest and most influential high-tech corporations, are not so much the humans who own them but the algorithms themselves.

One of the prime challenges for humanity in the age of AI will be to ensure that humans will continue to prosper and obtain a fair share of the resources produced by our shared human-AIA economy. The themes and policy proposals of the preceding sections of this policy note take on even greater urgency when viewed through this lens: Undoing technological redistribution and steering technological progress are even more important when they are about the distribution of resources between humans and artificial entities. Moreover, reducing the monopoly power of digital superstars gains extra importance when it is about maintaining the consumption share of humans in our common economy.

In spite of all these measures, human labor may well become irrelevant in the labor market in coming decades (Korinek and Stiglitz, 2018, 2019). Satisfying the basic needs of us humans would then require income from sources other than labor, whether they be labeled a social dividend, an allocation of subsistence income, or a universal basic income. The political difficulty of direct handouts can be reduced by providing many of the services that we humans rely on, such as healthcare and education, for free. Furthermore, some may view it desirable to subsidize humans to perform tasks that provide meaning, even though they are wasteful and redundant from an economic perspective. The stark alternative would be to let Malthusian forces play out, which would lead to large unnecessary suffering in a world of growing abundance.

Anton Korinek *is an Associate Professor of Economics at the University of Virginia. Contact:* *anton@korinek.com*

Endnotes

* I would like to thank Susan Helper, Rebecca Henderson, Ioana Marinescu and Dani Rodrik for their thoughtful comments and Hasan Toprak for excellent research assistance. Special thanks go to João Oliveira for sharing data to calculate the human capital share of income.

References

Acemoglu, Daron and Pascual Restrepo (2019), "The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand," Working Paper.

Agrawal et al. (2019), *The Economics of Artificial Intelligence*, NBER and University of Chicago Press.

Autor, David (2015), Why are there still so many jobs? The History and Future of Workplace Automation, *Journal of Economic Perspectives* 29(3).

Brinca, Pedro, João B. Duarte and João G. Oliveira (2019), "Investment-specific Technological Change, Taxation and Inequality in the U.S." *MPRA Paper 91463*, University Library of Munich, Germany.

Korinek, Anton (2018), "The Rise of Artificially Intelligent Agents," Working Paper.

Korinek, Anton and Ding Xuan Ng (2018), "Digitization and the Macroeconomics of Superstars," Working Paper.

Korinek, Anton and Joseph E. Stiglitz (2019), "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," forthcoming in Agrawal et al.: *The Economics of Artificial Intelligence*, NBER and University of Chicago Press.

Korinek, Anton and Joseph E. Stiglitz (2018), "Worker-Replacing Technological Progress," Working Paper.