Designing Policy to Steer Progress in Artificial Intelligence[1]

by Anton Korinek

University of Virginia, Department of Economics and Darden School of Business

Abstract

Artificial intelligence has recently led to widespread anxiety about potential future job losses. This memo takes as a premise that it is desirable to offer well-paying jobs to all able-bodied workers – either because jobs are important to provide meaning or because it is politically too costly to provide income in alternative ways such as via transfers. I therefore investigate how to steer technological progress so as to increase labor demand and create better-paying jobs, especially for lesser-educated workers.

Introduction

Artificial intelligence has advanced rapidly in recent years and has led to widespread anxiety that a growing number of human jobs will soon be made redundant. Over the next decade or two, Frey and Osborne (2013) predict that 47% of US jobs could be automated. The story of technological progress has always involved job losses. However, despite all the job losses along the way, the Industrial Revolution as a whole has substantially increased labor demand and has thus led to large increases in wage levels and overall welfare over the past 250 years. What is different about recent waves of automation is that they appear to be strongly labor-saving, i.e. they reduce equilibrium wage levels even though economic output is rising. (See e.g. the evidence on industrial robots reported by Acemoglu and Restrepo, 2017).

Some technologists go even further and worry that within the next several decades, artificial intelligence will be able to perform substantially all human jobs more cheaply than the subsistence cost of human labor and threaten to make human labor economically redundant (see e.g. Kurzweil, 2005; Bostrom, 2014). Such alarmist predictions are of course speculative and subject to considerable uncertainty. Nonetheless, they suggest that it may be a good idea to think carefully about the topic.

This memo proposes to respond to these challenges by actively steering technological progress towards inventions that augment human labor rather than replacing it. Technological progress does not just happen by itself. It would be misplaced to succumb to techno-fatalism and view our fate as pre-determined by blind technological forces and market forces that are beyond our control. Instead, our

[1] This memo draws on work-in-progress on "Steering Progress in Artificial Intelligence" as well as my papers on "Integrating Ethical Values and Economic Value to Steer Progress in Artificial Intelligence" (2020), "Preparing for the (Non-Existent?) Future of Work" (2020) and "Artificial Intelligence and Its Implications for Income Distribution and Unemployment" (2019).

future is shaped jointly by the technological innovations that we humans create, by the social and economic institutions that we collectively design, and by the ethical values that guide it all. We as a society have the power to actively steer the path of technological progress in AI through individual and collective action so as to shape the future that we want to live in.

The Value of Jobs

Jobs serve two crucial functions in our society that are at the heart of human existence: Jobs are not only the main source of income for the majority of people, but they are also one of the main activities on which people spend their time. As a result, our jobs play a central role in shaping our identity and social connections. This unholy alliance of economic necessity and shaping our identity implies that the possibility of a future in which a growing number of people can no longer find gainful employment would pose not only significant economic challenges but also shake the broader foundations of our society.Let us unpack the two components of jobs a little more in the following.

Jobs as a Source of Income

First, jobs serve as the main source of income for the majority of the population. Most people earn their living from paid work, either from their own present-day work, from a family member's work or, in the case of retirees, from savings and social security claims that they have accumulated by their past work.

Substituting for people's labor income would pose both fiscal and political challenges. For example, one of the most popular policy solutions that has been suggested to deal with widespread job losses in the future – a Universal Basic Income (UBI) that pays every adult a sufficient amount to meet basic subsistence needs – is rather expensive: For example, Kearney and Mogstad (2019) calculate that even a basic UBI that pays $10,000 to every U.S. adult a year would cost about $2.5 trillion per year, which would represent well more than half the current federal budget.

Jobs as a Source of Identity, Connections, Status, and Meaning

Jobs not only serve as a source of income, but they also involve many non-economic factors. Our jobs structure our time, shape many of our social connections and our social status, and ultimately also provide us with a sense of identity, self-esteem and meaning. People who lose their jobs simultaneously lose all these factors, which exacerbate the economic losses that they experience. Studies have shown that people who lose their jobs experience declines in both physical and mental health, social withdrawal, family disruptions, and even negative implications on their children's educational attainment.

Put differently, jobs serve a variety of non-economic functions in our society that provide benefits that go far beyond the pure economic output that they generate. Our society is organized around jobs and, at least in the short- to medium-run, it seems difficult to find a wholesale replacement for the important role in our society played by jobs.

Varieties of technological progress

Technological progress by definition increases economic output. However, how the fruits of the progress are shared – and whether there are perhaps losers from progress – depends on the specific nature of the innovations involved. Looking at relative gains, economists call technological progress labor-biased if labor benefits proportionately more than capital, and capital-biased otherwise. If labor and capital

benefit in equal proportions, then economists call it neutral technological progress. Broadly speaking, technological progress since the Industrial Revolution has been roughly neutral, although there was some bias towards capital in recent decades.

Looking at absolute gains and losses, we call technological progress labor-using if it increases wages and labor-saving if it reduces wages. Economists used to think that progress is in practice always labor-using, i.e. that it increases wages even when it was capital-biased and capital actually gained more than labor. However, looking at recent data, the realization has set in that some forms of technological progress are labor-saving (see e.g. Acemoglu and Restrepo, 2017).

Worries about Technology Eating Jobs

Whereas much of the popular press speaks of worries about artificial intelligence leading to job losses, which captures the subjective experience of workers who lose their jobs due to automation, the main worry of economists is that artificial intelligence will be strongly labor-saving. Job losses are painful for the individual, but when the lost jobs are replaced by better jobs, the pain is temporary. By contrast, labor-saving technological progress implies that lost jobs will be replaced by worse jobs that pay lower wages and/or offer less attractive working conditions.

If an increasing number of human jobs could be performed by machines and AI systems at some point in the future, the market wages of humans would be determined by how cheaply machines can perform their jobs. This would result in ever-declining wages as machines continue to become more efficient. At some point, human wages may decline below subsistence levels, i.e. human workers could no longer survive solely from their wage income. When the popular press simplistically speaks of a future without jobs, what economists really worry about is a future without a living wage. In such a future, workers may still be able perform jobs, but they could no longer survive from their wages alone.

To assess whether we are navigating towards a job-less future, the main economic indicator to look at is therefore not the number of human jobs but the level of wages. At the time of writing, it does not look like we need to worry about a job-less future on the immediate horizon. Real-wages of US workers have been growing at a healthy pace in recent years. However, there are some warning signs on the horizon if we look at longer periods. Over the past four decades, median real wages of US workers have risen by less than 10%, whereas real GDP/capita has approximately doubled. Real wages of non-supervisory workers have in fact declined. Although there are many forces at play that have contributed to this development (including globalization and institutional changes), this provides a warning sign that the economy no longer values human labor the way it used to and that recent technological progress has left the average worker behind.

The main theme of this memo is that it doesn't have to be that way. We as a society have a choice. We should not view our fate as pre-determined by blind technological forces and market forces that are beyond our control. Instead, we have the power to actively steer technological progress towards inventions that are labor-using instead of labor-saving in order to ensure that workers obtain not only the economic benefits of jobs but also the identity, connections, status and meaning that jobs provide to members of our society.

In the following, we will first discuss how to identify which innovations are labor-using versus labor-saving; then we will investigate policy measures to steer progress towards labor-using innovations.

Identifying Labor-Using Innovations

One of the main challenges in steering technological progress is to identify what the impact of a given innovation on labor market will be. A technological innovation is by definition not known before it is invented, and this generally makes it difficult to judge its economic effects.

Conceptually speaking, we expect the benefit of innovations for labor to be greater

(i) the larger the market for the product produced,
(ii) the more complementary the innovation is to labor,
(iii) the greater the labor share of the product produced,
(iv) the greater the demand elasticity for the product, i.e. the more demand for the product goes up when innovation makes the good cheaper,

Examples of Labor-Using Innovations

To make the description above more concrete, let me list two specific examples of labor-using innovations in the space of artificial intelligence:

1) Intelligent assistants are AI-powered devices that assist human workers and increase their productivity by complementing their cognitive capabilities. A specific application for which such assistants have recently been proposed are Augmented-Reality devices that help upskill lesser-skilled workers by providing them with specific instructions on how to perform cognitively intensive jobs. For example, such devices can assist factory workers perform complicated workflows that would otherwise require significant training. Another application are AI systems that provide call center workers with additional information about the callers in question, e.g. by analyzing the emotional content of voices. Even systems such as Google Maps can be interpreted as intelligent assistants that provide driving instructions to human drivers and thereby allow them to navigate routes in areas that they are not familiar with. One potential downside of intelligent assistants is that they may actually lower the skill levels of workers because they make them dependent on the assistants, and that they turn human workers that used to think for themselves into “robots” that mechanically follow the instructions given by the assistant.

2) Platforms that match labor services are another example of labor-using innovations. A number of high-tech corporations specialize in matching demand and supply for labor in the economy. Specific examples include Uber or Lyft, which match demand and supply for drivers, or Etsy, which matches demand and supply for artisan goods. Although there are justified concerns about the specifics of the jobs created e.g. by ride-sharing companies, these concerns can be addressed by appropriate regulation. In any case, platform companies that match labor have a labor-using effect because they increase overall demand for labor in the economy.

Policy Measures to Steer Progress

My proposal to steer progress in AI to maximize the positive impact on average workers is applicable in a number of domains:

First, many innovators and entrepreneurs in the technology sector are eager to maximize the positive impact of their developments on society. At present, however, the impact of technological progress on labor markets and income distribution is all too often an afterthought for innovators. Publicly-spirited

innovators will find it useful to be reminded of and obtain better guidance on the likely impact of their inventions. If the world's most creative innovators put their minds to it, they can play an important positive role in guiding progress in a direction that is beneficial for the average worker. Furthermore, innovators are perhaps also best-suited to predict the potential implications of their innovations and make better-informed decisions on what innovations to pursue to further the interests of workers.

Second, works councils that have a say in which types of investments and innovations to pursue in their companies are also well-suited in judging the effects of specific innovations on workers. If they have the right to participate in the decision-making process, they will steer technological progress in a direction that is positive for their members. This is the precise opposite of the efforts of some corporations to make their workers as replaceable as possible in order to reduce workers' bargaining power.

Third, a significant part of AI research is either conducted or sponsored by government. Although this type of research is funded by the tax dollars of all workers, the government typically pays little attention to how the resulting innovations affect the livelihood of workers. A natural public policy is thus to evaluate the likely effect of innovations on labor markets when determining what type of research the government should pursue.

Fourth, whether intentionally or unintentionally, our tax system plays an important in affecting the direction of technological progress: at present, labor is the most highly-taxed factor in our economic system, creating strong incentives for labor-saving innovation. One of the most natural public policy measures to steer progress in a direction that augments human labor is to reduce the burden of taxation on labor or, perhaps at some point, subsidize human labor.

Last but not least, it may also be desirable to provide explicit economic incentives for innovative efforts to benefit human labor. Provided that it is possible to identify innovations that are definitely labor-using, it may be desirable to directly subsidize them.

Selected References:

Acemoglu, Daron and Pascual Restrepo (2017), "Robots and Jobs: Evidence from US Labor Markets," NBER Working Paper 23285.

Acemoglu, Daron and Pascual Restrepo (2019), "The Wrong Kind of AI? Artificial Intelligence and the Future of Labor Demand", NBER Working Paper 25682.

Brand, Jennie E (2015), "The Far-Reaching Impact of Job Loss and Unemployment," *Annual Review of Sociology* 41, pp. 359-375.

Kearney, Melissa and Magne Mogstad (2019), "Universal Basic Income (UBI) as a Policy Response to Current Challenges," working paper, Aspen Economic Strategy Group.

Korinek, Anton (2020), "Integrating Ethical Values and Economic Value to Steer Progress in Artificial Intelligence," forthcoming in Dubber, Markus, Frank Pasquale and Sunit Das (eds.), *Oxford Handbook of Ethics of Artificial Intelligence*, Oxford University Press.

Korinek, Anton and Megan Juelfs (2020), Preparing for the (Non-Existent?) Future of Work, Technical note, Darden Center for Business in Society.

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