



Robot holds newspaper during 2016 World Economic Forum, Davos, Switzerland.

Robots, Growth, and Inequality

Andrew Berg, Edward F. Buffie, and Luis-Felipe Zanna

The robot revolution could have profound negative implications for equality

SOME say the world is entering a “second machine age.” Every week we read about a new application of artificial intelligence, so-called deep learning, and robotic technology. Automated delivery trucks, electronic teaching and scheduling assistants, computers that replace paralegals, and self-driving cars are just a few. Some seem to approach the “robot” envisioned by Czech science fiction writer Karel Čapek, who coined the term in 1921 to describe an intelligent machine essentially indistinguishable from a human.

No one knows where this technology is headed. Robert Gordon argues that economically meaningful technological change—and productivity growth in the United States—

has slowed since the 1970s, except for a decade-long tech boom ending in 2004 (see the June 2016 *F&D*). But when it comes to intelligent robots, we may be in the early stages of a revolution, and economists should think hard about what it means for economic growth and income distribution.

Competing narratives

Two narratives have emerged in the economic literature on technology, growth, and distribution. One says that technological advances raise productivity and thus output per person. Despite some transitional costs as particular jobs become obsolete, the overall effect is a higher standard of living. The history of this debate since at least the 19th century seems to

yield a decisive victory for technological optimists. The average American worker in 2015 worked roughly 17 weeks to live at the annual income level of the average worker in 1915—and technology was a huge part of that progress (Autor, 2014).

This optimistic narrative points to the many ways that technology does much more than displace workers. It makes workers more productive and raises demand for their services—for example, mapping software makes taxi (and now Lyft and Uber) drivers more efficient. And rising incomes generate demand for all sorts of outputs and hence labor. A wave of fear about the implications of computerization for jobs surged in the United States in the 1950s and early 1960s, but subsequent decades of strong productivity growth and rising standards of living saw roughly stable unemployment and rising employment.

The other, more pessimistic, narrative pays more attention to the losers (see, for example, Sachs and Kotlikoff, 2012; Ford, 2015; Freeman, 2015). Some of the increased inequality in many advanced economies in recent decades may result from technological pressure. The computer revolution has reduced relative demand in developed economies for jobs involving routinized work (physical or mental)—think bookkeeper or factory line worker. Because computers combined with a smaller number of—generally more skilled—workers have been able to produce the goods previously associated with these jobs, relative wages for people with fewer skills have fallen in many countries.

Will robots be different?

Where might intelligent robots fit in? For a bird's-eye view of this question, we designed an economic model that assumes robots to be a different sort of capital, one that is a close substitute for human workers. Macroeconomists usually think of production as resulting from the combination of physical capital stock (comprising machines and structures, both public and private) and labor. But thinking of robots as a new type of physical capital, one that in effect adds to the stock of available (human) labor, is surprisingly instructive. Production will still require buildings and roads, for example, but now people and robots can work with this traditional capital.

So what happens when this robot capital gets productive enough to be useful? If we assume that robots are almost perfect substitutes for human labor, the good news is that output per person rises. The bad news is that inequality worsens, for several reasons. First, robots increase the supply of total effective (workers plus robots) labor, which drives down wages in a market-driven economy. Second, because it is now profitable to invest in robots, there is a shift away from investment in traditional capital, such as buildings and conventional machinery. This further lowers the demand for those who work with that traditional capital.

But this is just the beginning. Both the good and bad news intensify over time. As the stock of robots increases, so does the return on traditional capital (warehouses are more useful with robot shelf stockers). Eventually, therefore, traditional investment picks up too. This in turn keeps robots productive, even as the stock of robots continues to grow. Over time,

the two types of capital grow together until they increasingly dominate the entire economy. All this traditional and robot capital, with diminishing help from labor, produces more and more output. And robots are not expected to consume, just produce (though the science fiction literature is ambiguous about this!). So there is more and more output to be shared among actual people.

However, wages fall, not just in relative terms but absolutely, even as output grows.

This may sound odd, or even paradoxical. Some economists talk about the fallacy of technology fearmongers' failure to realize that markets will clear: demand will rise to meet the

If robots are almost perfect substitutes for human labor, inequality worsens.

higher supply of goods produced by the better technology, and workers will find new jobs. There is no such fallacy here: in our simple model economy, we assume away unemployment and other complications: wages adjust to clear the labor market.

So how can we explain the fall in wages coinciding with the growing output? To put it another way, who buys all the higher output? The owners of capital do. In the short run, higher investment more than counterbalances any temporary decline in consumption. In the long run, the share of capital owners in the growing pie—and their consumption spending—is itself growing. With falling wages and rising capital stocks, (human) labor become a smaller and smaller part of the economy. (In the limiting case of perfect substitutability, the wage share goes to zero.) Thomas Piketty has reminded us that the capital share is a basic determinant of income distribution. Capital is already much more unevenly distributed than income in all countries. The introduction of robots would drive up the capital share indefinitely, so the income distribution would tend to grow ever more uneven.

An economic robot “singularity”?

Remarkably, this process of self-sustained purely investment- (robot plus traditional) driven growth can take off even with a very small increase in robots' efficiency, as long as this increase makes robots competitive with labor. This tiny efficiency boost thus leads to a sort of economic “singularity,” in which capital takes over the entire economy to the exclusion of labor. It is reminiscent of the hypothesis of “technological singularity” publicized by Raymond Kurzweil (2005) in which intelligent machines become so smart that they can program themselves, triggering explosive further growth of machine intelligence. Ours is an economic, not a technological, singularity, however. We are considering how a small jump in the level of robot efficiency could trigger self-sustaining capital accumulation whereby robots take over the economy, not self-sustaining growth in robot intelligence.



Humanoid robot Pepper takes selfie during app contest, Tokyo, Japan.

So far, we've assumed nearly perfect substitutability between robots and workers along with a small increase in robot efficiency. These are robots of the sort featured in the Hollywood movie *Terminator 2: Judgment Day*—such perfect substitutes for humans that they are indistinguishable. Another plausible scenario departs from both these assumptions. It is more realistic, at least for now, to assume that robots and human labor are close *but not perfect* substitutes, that people bring a spark of creativity or a critical human touch. At the same time, like some technologists, we project that robot productivity increases not just a little but dramatically over a span of a couple of decades.

With these assumptions, we recover a bit of the economist's typical optimism. The forces mentioned before are still at play: robot capital tends to replace workers and drive down wages, and at first the diversion of investment into robots dries up the supplies of traditional capital that help raise wages. The difference, though, is that humans' special talents become increasingly valuable and productive as they combine with this gradually accumulating traditional and robot capital. Eventually, this increase in labor productivity outweighs the fact that the robots are replacing humans, and wages (as well as output) rise.

But there are two problems. First, "eventually" can be a long time coming. Exactly how long depends on how easy it is to substitute robots for human labor, and how quickly savings and investment respond to rates of return. According to our baseline calibration, it takes 20 years for the productivity effect to outweigh the substitution effect and drive up wages. Second, capital will still likely greatly increase its role in the economy. It will not completely take over as it does in the singularity case, but it will take a higher share of income, even in the long run when wages are above the pre-robot-era level. Thus, inequality will be worse, possibly dramatically so.

People are different

Readers may be thinking that these scary scenarios will not apply to them, because their jobs as, say, economists or journalists cannot be performed by robots. In our model, we started

with labor and robots as perfect substitutes, then introduced the notion that they may be close but not perfectly the same in production. A further important complication is that not all labor is the same. And indeed, it is plausible that even sophisticated machines combined with advanced artificial intelligence will not replace humans for all jobs. In movies the range of jobs to be replaced is quite broad, from robot hunter (*Blade Runner*) to doctor (*Alien*). And robots have at least taken a stab at replacing teaching assistants and even journalists. Massive online courses may threaten even professors. But in real life, many jobs do seem safe, at least for now.

In our model, we therefore next divide all workers into two categories, which we call "skilled" and "unskilled." By skilled we mean that they are *not* close substitutes for robots; rather, robots may increase their productivity. By unskilled we mean that they are very close substitutes. Thus, our skilled workers may not be the traditionally highly educated; they may be those with creativity or empathy, which is particularly hard for future robots to match. We assume, following Frey and Osborne (2013), that about half of the labor force can be replaced by robots and is thus "unskilled." What happens when robot technology becomes cheaper? As before, output per person grows. And the share of overall capital (robots plus traditional) rises. Now, though, there is an additional effect: the wages of skilled workers rise relative to those of the unskilled—and absolutely. Why? Because these workers are more productive when combined with robots. Imagine, for example, the greater productivity of a designer who now commands an army of robots. Meanwhile, the wages of the unskilled collapse, both in relative and in absolute terms, even over the long run.

Inequality now increases for two fundamental reasons. As before, capital receives a greater share of total income. In addition, wage inequality worsens dramatically. Productivity and real wages paid to skilled labor increase steadily, but low-skilled workers wage a lonely battle against the robots and lose badly. The numbers depend on a few key parameters, such as the degree of complementarity between skilled work-

ers and robots, but the rough magnitude of the outcome follows from the simple assumptions we have laid out. We find that over a period of 50 miserable years, the real wage for low-skilled labor decreases 40 percent, and the group's share in national income drops from 35 percent to 11 percent in our baseline calibration.

Humans' special talents become increasingly valuable and productive as they combine with robot capital.

So far, we have been thinking of a large developed economy, like the United States. And this seems natural given that such countries tend to be more advanced technologically. However, a robot age could also affect the international distribution of output. For example, if the unskilled labor replaced by robots resembles the workforce of developing economies, it could lower those countries' relative wages.

Who will own the robots?

These stories are not destiny. First, we are mainly speculating about the outcome of emerging technological trends, not analyzing existing data. Recent innovations we have in mind have not (yet) shown up in productivity or growth statistics in developed economies; productivity growth has in fact been low in recent years. And technology does not seem to be the culprit for the rise in inequality in many countries. In most advanced economies growth in the relative wages of skilled workers has been smaller than in the United States, even in advanced economies presumably facing similar technological changes. As Piketty and his coauthors have famously emphasized, much of the increase in inequality in recent decades is concentrated in a very small fraction of the population, and technology does not seem to be the main story. But the rising inequality observed in so many parts of the world over recent decades—and perhaps even some of the political instability and populism in the news—underscores the risks and raises the stakes. And it is ominous that the labor share of income in the United States seems to have been falling since the turn of the century, after decades of rough stability (Freeman, 2015).

Science fiction writer Isaac Asimov's famous three "laws of robotics" were designed to protect people from physical harm by robots. According to the first law, "A robot may not injure a human being or, through inaction, allow a human being to come to harm." Such guidance may be fine for designers of individual robots, but it would do little to manage the economy-wide consequences we discuss here. Our little model shows that, even in a smoothly functioning market economy, robots may be profitable for owners of capital and may raise average per capita income, but the result would not be the kind of society most of us would want to live in. The case for a public policy response is strong.

In all these scenarios, there are jobs for people who want to work. The problem is that most of the income goes to owners of capital and to skilled workers who cannot easily be replaced by robots. The rest get low wages and a shrinking share of the pie. This points to the importance of education that promotes the sort of creativity and skills that will complement—not be replaced by—intelligent machines. Such investment in human capital could raise average wages and lower inequality. But even so, the introduction of robots may depress average wages for a long time, and the capital share will rise.

In trying to keep things as simple as possible, we have ignored many of the obligations such a society would face. These could include ensuring sufficient aggregate demand when buying power is increasingly concentrated, addressing the social and political challenges associated with such low wages and high inequality, and dealing with the implications of lower wages when it comes to workers' ability to pay for health care and education and invest in their children.

We have implicitly assumed so far that income from capital remains highly unequally distributed. But the increase in overall output per person implies that everyone could be better off if income from capital is redistributed. The advantages of a basic income financed by capital taxation become obvious. Of course, globalization and technological innovation have made it, if anything, easier for capital to flee taxation in recent decades. Our analysis thus adds urgency to the question "Who will own the robots?" ■

Andrew Berg is a Deputy Director in the IMF's Institute for Capacity Development, Edward F. Buffie is a Professor of Economics at Indiana University Bloomington, and Luis-Felipe Zanna is a Senior Economist in the IMF's Research Department.

References:

- Autor, David, 2014, "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, Vol. 29, No. 3, pp. 3–30.
- Čapek, Karel, 1921, R.U.R. (Rossum's Universal Robots) (New York: Penguin).
- Ford, Martin, 2015, *The Rise of the Robots* (New York: Basic Books).
- Freeman, Richard B., 2015, "Who Owns the Robots Rules the World," *IZA World of Labor*, May.
- Frey, Carl Benedikt, and Michael A. Osborne, 2013, "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Oxford University paper* (Oxford, United Kingdom).
- Gordon, Robert, 2016, *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War* (Princeton, New Jersey: Princeton University Press).
- Kurzweil, Raymond, 2005, *The Singularity Is Near: When Humans Transcend Biology* (New York: Viking).
- Sachs, Jeffrey D., and Laurence Kotlikoff, 2012, "Smart Machines and Long-Term Misery," *NBER Working Paper 18629* (Cambridge, Massachusetts: National Bureau of Economic Research).