



project-syndicate.org Sami Mahroum

The AI Debate We Need by Sami Mahroum

Feb 16, 2018

Rapid advances in artificial intelligence and related technologies have contributed to fears of widespread job losses and social disruptions in the coming years, giving a sense of urgency to debates about the future of work. But such discussions, though surely worth having, only scratch the surface of what an AI society might look like.

BARCELONA – One can hardly go a day without hearing about a new study describing the far-reaching implications of advances in artificial intelligence. According to countless consultancies, think tanks, and Silicon Valley celebrities, AI applications are poised to change our lives in ways we can scarcely imagine.

The Year Ahead 2018

The world's leading thinkers and policymakers examine what's come apart in the past year, and anticipate what will define the year ahead.

Order now

The biggest change concerns employment. There is widespread speculation about how many jobs will soon fall victim to automation, but most forecasters agree that it will be in the millions. And it is not just blue-collar jobs that are at stake. So, too, are high-skilled white-collar professions, including law, accounting, and medicine. Entire industries could be disrupted or decimated, and traditional institutions such as universities might have to downsize or close. Such concerns are understandable. In the current political economy, jobs are the main vehicle for wealth creation and income distribution. When people have jobs, they have the means to consume, which drives production forward. It is not surprising that debates about AI would center on the prospect of mass unemployment, and on the forms of compensation that could become necessary in the future.

But, to understand better what AI will mean for our shared economic future, we should look past the headlines. We can start with insights from *Project Syndicate* commentators, who assess AI's economic implications by situating the current technological revolution in a larger

historical context. Their analyses suggest that AI will indeed reshape employment across advanced and developing economies alike, but also that the future of work will be but one small part of a much larger story.

From Each AI According to Its Abilities...

For Nobel laureate economist [Christopher Pissarides and Jacques Bughin](#) of the McKinsey Global Institute, the AI revolution need not “conjure gloom-and-doom scenarios about the future of work,” so long as governments rise to the challenge of equipping workers “with the right skills” to prepare them for future market needs. Pissarides and Bughin remind us that job displacement from new technologies is nothing new, and often comes in waves. “But throughout that process,” they note, “productivity gains have been reinvested to create new innovations, jobs, and industries, driving economic growth as older, less productive jobs are replaced with more advanced occupations.”

SAP CEO [Bill McDermott](#) is similarly optimistic, and sees “nothing to be gained from fearing a dystopian future that we have the power to prevent.” Rather than rendering humans obsolete, McDermott believes that AI applications could liberate millions of people from “the dangerous and repetitive tasks often associated with manual labor.” And he points to the introduction of “collaborative robots” to show that “partnership, not rivalry” will define our future relationship with AI technologies across all sectors. But, as McDermott is quick to point out, this worker-machine dynamic will not come about on its own. “Those leading the change” must not lose sight of the “human element,” or the fact that “there are some things even the smartest machines will never do.”



But, as [Laura Tyson](#) of the University of California, Berkeley, warns, the design of new “smart machines” is less important than “the policies surrounding them.” Tyson notes that technological change has, in fact, already been displacing workers for three decades, accounting for an estimated 80% of the job losses in US manufacturing. In her view, we could be heading for a “‘good-jobless future,’ in which a growing number of workers can no longer earn a middle-class income, regardless of their education and skills.” To minimize that risk, she calls on policymakers in advanced economies to “focus on measures that help those who are displaced, such as [education and training programs](#), and income support and social safety nets, including wage insurance, lifetime retraining loans, and portable health and pension benefits.”

Alongside those welcoming or worrying about AI are others who consider current warnings to be premature. For example, Tyson’s University of California, Berkeley, colleague [J. Bradford DeLong](#) believes that, “it is profoundly unhelpful to stoke fears about robots, and to frame the issue as ‘artificial intelligence taking American jobs.’” Taking a long historical view, DeLong argues that there have been “relatively few cases in which technological progress, occurring within the context of a market economy, has directly impoverished unskilled workers.” Still, like Tyson, he notes that “workers must be educated and trained to use increasingly high-tech tools,” and that redistributive policies will be needed to “maintain a proper distribution of income.”

The Options on the Table

Containing income inequality is in fact one of the primary challenges of the digital age. One possible remedy is a tax on robots, an idea first [proposed](#) by Mady Delvaux of the European Parliament and later endorsed by Microsoft founder [Bill Gates](#). Nobel laureate economist [Robert Shiller](#) observes that while the idea has drawn derision in many circles, it deserves an airing, because there are undeniable “externalities to robotization that justify some government intervention.” Moreover, there aren’t any obvious alternatives, given that “a more progressive income tax and a [‘basic income’](#)” lack “widespread popular support.”

But [Yanis Varoufakis](#) of the University of Athens sees another solution: “a [universal basic dividend](#) (UBD), financed from the returns on all capital.” Under Varoufakis’s scheme, the pace of automation and rising corporate profitability would pose no threat to social stability, because society itself would become “a shareholder in every corporation, and the dividends [would be] distributed evenly to all citizens.” At a minimum, Varoufakis contends, a UBD would help citizens recoup or replace some of the income lost to automation.

Similarly, [Kaushik Basu](#) of Cornell University thinks there should be a larger focus “on expanding profit-sharing arrangements, without stifling or centralizing market incentives that are crucial to drive growth.” Practically speaking, managing the rise of new tech monopolies that enjoy unjustifiable “returns to scale” would require giving “all of a country’s residents the

right to a certain share of the economy's profits." At the same time, it will mean replacing "traditional anti-monopoly laws with legislation mandating a wider dispersal of shareholding within each company."

Another option, notes [Stephen Groff](#) of the Asian Development Bank, is to direct workers toward fields that will not necessarily fall prey to automation. For example, "governments should offer subsidies or tax incentives to companies that invest in the skills that humans master better than machines, such as communication and negotiation." Another idea, notes [Kemal Derviş](#) of the Brookings Institution, is a "job mortgage," whereby firms "with a future need for certain skills would become a kind of sponsor, involving potential future job offers, to a person willing to acquire those skills."

And at a more fundamental level, argues [Andrew Wachtel](#), the president of the American University of Central Asia, we should be preparing people for an AI future by teaching "skills that make humans human." The workers of tomorrow, he notes, "will require training in ethics, to help them navigate a world in which the value of human beings can no longer be taken for granted."

Stepping Off the Treadmill

And yet, as useful as these ideas are, they do not address a fundamental question of the digital age: Why do we still need jobs? After all, if AI technologies can deliver most of the goods and services that we need at less cost, why should we spend our precious time laboring? The impulse to preserve traditional employment is an artifact of the industrial age, when the work-to-consume dynamic drove growth. But now that capital growth is outpacing job growth, that model is breaking down.

Capital, land, and labor were the three pillars of the industrial age. But digitalization and the so-called platform economy have devalorized land, and the AI revolution now threatens to render much labor obsolete. The question for a fully automated future, then, is whether jobs can be delinked from incomes, and incomes delinked from consumption. If not, then we could be headed for what [Robert Skidelsky](#) of Warwick University describes as "a world in which we are condemned to race with machines to produce ever-larger quantities of consumption goods."

Fortunately, the AI revolution holds out the promise of an alternative future. As [Adair Turner](#) of the Institute for New Economic Thinking points out, it is not hard to imagine "a world in which solar-powered robots, manufactured by robots and controlled by artificial intelligence systems, deliver most of the goods and services that support human welfare." At the same time, the social theorist Jeremy Rifkin, in [The Zero Marginal Cost Society](#), shows how shared platforms could produce countless new goods and services, and how new business models might emerge to monetize those platforms, all at no cost to consumers.

If this sounds farfetched, consider that it is already happening. Billions of people around the world now use platforms such as Facebook, WhatsApp, and Wikipedia for free. As DeLong [notes](#), “More than ever before, we are producing commodities that contribute to social welfare through use value rather than market value.” And people are spending increasingly more time “interacting with information-technology systems where the revenue flow is, at most, a tiny trickle tied to ancillary advertising.”

As it advances, AI could allow us to consume ever more products and services from an expanding “freemium” economy based on network effects and “collective intelligence,” not unlike an open-source community. At the same time, agents in a parallel premium economy will continue to mine AI-based systems to extract new value. In an advanced AI economy, fewer people would hold traditional jobs, governments would collect less in taxes, and countries would have smaller GDPs; yet everyone would be better off, free to consume a widening range of goods that have been decoupled from income.

The End of Employment

In such a scenario, a job would become a luxury or hobby rather than a necessity. Those looking for more income would most likely have opportunities to do so through data mining, in the same way that cryptocurrency miners do today. But, because such income would be useful only for purchasing products and services that have resisted AI production, trading would be consigned to niche markets operated through blockchain networks. As [Maciej Kuziemski](#) of the University of Oxford puts it, AI will not just “change human life,” but will also alter “the boundaries and meaning of being human,” beginning with our self-conception as laboring beings.

Again, this may sound farfetched, or even utopian; but it is a more realistic depiction of the future than what one hears in current debates about preserving industrial-era economic frameworks. For example, plenty of people – not least rent-seeking owners of capital – already do not make a living from selling their labor. In an AI society, we could expect to see the protestant work ethic described by Max Weber gradually become an anachronism. Work would give way to higher forms of human activity, as the German philosopher Josef Pieper envisioned. “The modern person works in order to live, and lives in order to work,” Pieper [observed](#) more than 70 years ago. “But, for the ancient and medieval philosopher, this was a strange way to view the world; rather, they would say that we work in order to have leisure.” In an AI economy, individuals might earn “income” from their data when they partake in physical recreation; make “green” consumption choices; or share stories, pictures, or videos. All of these activities already reap rewards through various apps today. But Princeton University’s [Harold James](#) believes the replacement of work with new forms of leisure poses significant hazards. In particular, James worries that AI’s cooptation of most mental labor will

usher in a “stupid economy,” defined by atrophying cognitive skills, just as technologies that replaced manual labor gave rise to sedentary lifestyles and expanded waistlines.

In my view, however, there is no reason to think that the technologies of the future will not provide even more opportunities for people to live smarter and more creatively. After all, reaping the full benefits of AI will itself require acts of imagination. Moreover, millions of people will be freed up to perform social work for which robots are unsuited, such as caring for children, the sick, the elderly, and other vulnerable communities. And millions more will engage in leisure-work activities in the freemium economy, where data will be the new “natural resource” *par excellence*.

Making Worklessness Work

Still, realizing this vision of the future is far from guaranteed. It will require not just new economic models, but also new forms of governance and legal frameworks. For example, Kuziemski argues that, “Empowering all people in the age of AI will require each individual – not major companies – to own the data they create.” [Hernando de Soto](#) of the Institute of Liberty and Democracy adds the corollary that ensuring equal access to data for all people will be no less important.

Such imperatives highlight the fundamental ethical questions surrounding the AI revolution. Ultimately, the regulatory and institutional systems that we create to manage the new technologies will reflect our most deeply held values. But that means regulations and institutions might evolve differently from country to country. This worries [Guy Verhofstadt](#) of the Alliance of Liberals and Democrats for Europe Group (ALDE) in the European Parliament, who urges his fellow Europeans to start setting standards for AI now, before governments with fewer concerns about privacy and safety do so first.

With respect to safety, University of Connecticut philosopher [Susan Leigh Anderson](#) argues that machines should be permitted “to function autonomously only in areas where there is agreement among ethicists about what constitutes acceptable behavior.” More broadly, she cautions those developing ethical operating protocols for AI technologies that “ethics is a long-studied field within philosophy,” one that “goes far beyond laypersons’ intuitions.”

Underscoring that point, Princeton University’s [Peter Singer](#) lists various ethical dilemmas that are already confronting AI developers, and which have no clear solution. For example, he wonders whether driverless cars “should be programmed to swerve to avoid hitting a child running across the road, even if that will put their passengers at risk.” Singer warns against thinking of AI as merely a machine that can beat a human in chess or go. “It is one thing to unleash AI in the context of a game with specific rules and a clear goal,” he writes. “It is something very different to release AI into the real world, where the unpredictability of the environment may reveal a software error that has disastrous consequences.”

The potential for AI to provoke a backlash will be particularly acute in public services, where robots might manage our personal records or interact with children, the elderly, the sick, or socially marginalized groups. As [Simon Johnson and Jonathan Ruane](#) of MIT Sloan remind us, “what is simple for us is hard for even the most sophisticated AI; conversely, AI often can do easily what we regard as difficult.” The challenge, then, will be to determine – and not only on safety grounds – where and when AI should and should not be deployed.

Furthermore, democracies, in particular, will need to establish frameworks for holding those in charge of AI applications accountable. Given AI’s high-tech nature, governments will most likely have to rely on third-party designers and developers to administer public-service applications, which could pose risks to the democratic process. But the University of Oxford ethicist [Luciano Floridi](#) fears the opposite scenario, in which “AI is no longer controlled by a guild of technicians and managers,” and has been made “available to billions of people on their smartphones or some other device.”

A Broad Agenda

At the end of the day, policymakers setting a course for the future must focus on ensuring a smooth passage into an AI-enabled freemium economy, rather than trying to delay or sabotage the inevitable. They should follow the example of policy interventions in earlier periods of automation. As New York University’s [Nouriel Roubini](#) reminds us, “late nineteenth- and early twentieth-century leaders” sought to “minimize the worst features of industrialization.”

Accordingly, child labor was abolished, working hours were gradually reduced, and “a social safety net was put in place to protect vulnerable workers and stabilize the (often fragile) macroeconomy.”

A more recent success has been “green” policies that give rise to new business models. Such policies include feed-in tariffs, carbon credits, carbon trading, and Japan’s “Top Runner” program. When thinking about the freemium economy, governments should consider introducing automation offsets, whereby businesses that adopt labor-replacing technologies must also introduce a corresponding share of freemium goods and services into the market. More broadly, policy approaches to education, skills training, employment, and income distribution should all now assume a post-AI perspective. As Floridi notes, this will require us to question some of our most deeply held convictions. “A world where autonomous AI systems can predict and manipulate our choices,” he observes, “will force us to rethink the meaning of freedom.” Similarly, we will also have to rethink the meaning and purpose of education, skills, jobs, and wages.

Moreover, we will have to re-conceptualize economic value for a context in which most things are free, and spending of any kind is a luxury. We will have to decide on appropriate forms of capital ownership under such conditions. And we will have to create new incentives for people to contribute to society.

All of this will require new forms of proprietary rights, new modes of governance, and new business models. In other words, it will require an entirely new socioeconomic system, one that we will either start shaping or allow to shape us.