

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: The Economics of Artificial Intelligence: An Agenda

Volume Authors/Editors: Ajay Agrawal, Joshua Gans, and Avi Goldfarb, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-61333-8 (cloth); 978-0-226-61347-5 (electronic)

Volume URL: <http://www.nber.org/books/agra-1>

Conference Date: September 13–14, 2017

Publication Date: May 2019

Chapter Title: Public Policy in an AI Economy

Chapter Author(s): Austan Goolsbee

Chapter URL: <http://www.nber.org/chapters/c14030>

Chapter pages in book: (p. 309 – 316)

Public Policy in an AI Economy

Austan Goolsbee

11.1 Introduction

This conference has brought together a mix of technology and economics scholars to think broadly about the role of artificial intelligence (AI) in the economy, and this short chapter will present a few thoughts about the role of policy in a world where AI becomes ubiquitous.

Most of the public discussion about an AI-dominated economy has focused on robots and the future of work. Ruminations by public figures like Bill Gates, Stephen Hawking, and Elon Musk have stoked fears that robots will destroy our jobs (and, possibly, the world). Some of these same figures have called for various heterodox policy ideas, too, from moving to colonies in space to taxing the robots to providing a universal basic income (UBI) untethered to work.

As the research and comments in this volume suggest, economists have generally been less pessimistic when thinking about the role of AI on jobs. They often highlight the historical record of job creation despite job displacement, documented the way technological advances have eliminated jobs in some sectors but expanded jobs and increased wages in the economy overall, and highlighted the advantages that the new technologies will likely have in the future (some recent discussions include Autor 2015; Autor and Salomons 2018; Brynjolfsson and McAfee 2014; Mokyr 2014).

Austan Goolsbee is the Robert P. Gwinn Professor of Economics at the University of Chicago Booth School of Business and a research associate of the National Bureau of Economic Research.

I wish to thank the participants at the NBER Artificial Intelligence conference for helpful comments. For acknowledgments, sources of research support, and disclosure of the author's or authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c14030.ack>.

The pessimistic case has come more from technology/business sector. Perhaps seeing the advances in technology up close, they worry that the machines may soon be so good that they could replace almost anyone. One major study across many industries by the McKinsey Global Institute (2017) argues that 73 million jobs may be destroyed by automation by 2030 because of the rise of the new technologies.

In many ways, it is unfortunate that labor market policy has dominated our thinking about the AI economy. The main economic impact of AI is not about jobs or, at least, is about much more than just jobs. The main economic impact of these technologies will be how good they are. If the recent advances continue, AI has the potential to improve the quality of our products and our standard of living. If AI helps us diagnose medical problems better, improves our highway safety, gives us back hours of our day that were spent driving in traffic, or even just improves the quality of our selfies, these are direct consumer benefits. These raise our real incomes and the economic studies valuing the improvements from quality and from new products tend to show their value is often extremely high (see the discussions in the volume of Bresnahan and Gordon [1997] or the discussions over valuing “free” goods like Goolsbee and Klenow [2006] and Varian [2013]).

That is a different way of saying that if AI succeeds, it will raise our productivity and higher productivity makes us rich. It is not a negative. Indeed, if AI succeeded in the way some fear, it would mean the exact reversal of the main problem facing growth in the last decade or more that productivity growth has been too slow. Indeed, it would decisively refute one of the central tenets of secular stagnationist thinkers like Gordon (2016), who argue that low productivity growth is a semi-permanent condition for the advanced economies because of the scarcity of path breaking ideas. Would that AI could change that equation.

This chapter will consider a few disparate thoughts about policy in an AI-intensive economy (interpreting AI broadly to include a cluster of information technology-based productivity improvements beyond just conventional artificial intelligence or machine learning). It will consider the speed of adoption of the technology—the impact on the job market and the implications for inequality across people and across places, discuss the challenges of enacting a universal basic income as a response to widespread AI adoption, discuss pricing, privacy and competition policy, and conclude with the question of whether AI will improve policy making itself.

11.2 The Speed of Adoption: Implications for the Job Market and for Inequality

Taking the issue of job displacement first, the basic conclusion of the economists is that for the last hundred years there have been massive

amounts of job displacement, yet the structural unemployment rate has not seemed to rise, much less trend toward 100 percent. Over time, people adjust. They move. They get skills. The long-run impact of labor-saving technologies has overwhelming been positive for market economies. If the fear is that AI will replace low-skill jobs, it is a fact that tens or even hundreds of millions of low-skill jobs were displaced by technology in previous years in a process very similar to the one we describe today. If the fear is that AI is different this time around because it will begin to replace types of jobs that have never been automated before like higher-skill or white-collar jobs, the historical data indicate that those groups have been able to adjust to shocks and move to new sectors and new geographic areas easier than lower-skill workers have.

A critical issue is, of course, how fast the adjustment takes place/the speed of adoption of AI technology. The economy has proven quite capable of inventing new things for people to do over the long run. Obviously, if change happens all at once, the adjustment problem is worst. Spread out over time, however, the adjustment can be manageable. Take the much discussed case of autonomous cars. There were about 3.5 million truck, bus, and taxi drivers in 2015, and suppose that every one of them were lost due to advances in self-driving car technology. If this loss takes place over fifteen years, this would average a little over 19,000 per month, and compare that to the fact that in 2017 the Job Openings and Labor Turnover Survey (JOLTS) data show that the economy generated about 5.3 million jobs per month (with 5.1 million separations per month). The complete elimination of every job in the sector would increase the separation rate by less than four-tenths of a percent. It would force drivers into new sectors and be disruptive to their livelihoods. But as a macroeconomic phenomenon, the impact would be small. If that loss happened in two years, the impact would be quite significant. So it is worth considering what influences the speed of adoption and, certainly, a key determinant will be how good the AI actually is compared to people. But, many analysts seem to view that as the *only* thing that will determine adoption rates. It is worth considering at least two other factors: prices and adjustment costs.

First, many of these AI innovations involve significant capital outlays up front and that alone may slow their adoption for some time. Ride-share drivers, for example, by some measures can barely cover the cost of operating their cars (including depreciation, fuel, maintenance, and insurance) at the price of cars now. AI-enabled autonomous vehicles are likely to cost substantially more per car than conventional cars when they become available to the public. Will companies be willing to incur large upfront costs to bypass paying drivers? It really depends on prices that we do not yet know.

Second, “better” does not always mean faster adoption. Economists have shown automated stock picking through index funds superior to active management for decades, yet people still hold trillions in inefficient,

high-fee funds. Millions of people have mortgages with higher than market interest rates that they do not refinance, cell phone data plans that do not match their usage, and so on. There are tens of millions of people that do not use the internet. Inertia is a powerful force slowing the adoption of technology products and is certainly worth remembering if we want to predict something like how fast people will give up common behaviors like driving for themselves.

Third, in an important sense, we know that AI can only be as good as its training sample and there are some very different types of customers in the country that may make the AI quality improvements much more fitting for certain types of customers than others. Microsoft created an AI program to learn from Twitter and see if it could create content that people would think was written by a human. They started it in the United States and had to shut it down almost immediately because it became so abusive and offensive. It mirrored what it saw online. Running the same program in China, where Twitter is heavily censored, it has performed well and not turned abusive. The attributes of the product and the “quality” of the product depend on how relevant the training sample is to that customer.

This is likely to influence the adoption rate of the AI technologies in different places. Again, think of the autonomous cars. Will we gather loads of information about driving in urban areas and on highways or in Silicon valley from the early adopters, tailor the product to their needs, but then find that it does not work as well for dirt roads or rural places or places without Bay Area weather?

Heterogeneous demand is the hobgoblin of the AI mind. Groups that differ most from the training sample will likely be the slowest to adopt the technology, in part, because it will be the least helpful to them. That may lead to another manifestation of the digital divide. In this sense, the rise of AI technologies is likely to make the problem of income and of geographic inequality even worse. To the extent that new AI technologies are expensive and tailored toward the training sample of adopters, it will be like having lower inflation and greater consumer surplus going to those groups (for discussions about differences in prices and innovation across income groups or for online buyers versus offline buyers, see Jaravel [2017] or Goolsbee and Klenow [2018]).

Government policy will face the potential of divisions along red state/blue state or high-education/low-education locations or high-income/low-income neighborhoods even more than it does today.

11.3 Challenges for Universal Basic Income as a Response to Job Market Displacement

Now suppose that the arguments above prove wrong. Nothing slows the speed of AI adoption and there is mass job displacement in a short time.

There has been a rising call among the believers in that scenario for universal basic income policy. Closely tied to the old Milton Friedman notion of a negative income tax, the UBI would grant some minimal level of income to people regardless of employment status as a new form of safety net and anyone could then work beyond that UBI level to earn more. In the purest libertarian concept, this UBI would replace the existing collection of safety net programs. The advantage of the UBI would be that people could survive in a world with few human jobs and alleviate poverty in a relatively efficient manner and without destroying all incentives in the private economy. It seeks to separate the notion of “making a living” from having a job. There are some small-scale experiments with the UBI in a few countries like Finland and New Zealand or funded by private individuals in the United States. There are a number of challenges associated with negative income taxes and UBIs as a policy solution to widespread AI adoption.

First, if you accept the economists’ basic labor supply model (that people value leisure and so generally need to be paid to work) then there are likely to be some sizable number of people who are working only because they absolutely have to. In a world where AI-induced unemployment is already high, separating work and income might be an advantage. In a world like the one we are in now, though, offering a basic income will likely cause a sizable drop in labor market participation by low-wage earners. To the extent that nonparticipation in exactly that segment of the labor force is already viewed as a problem, the UBI would likely make things worse and risk angering the broader public.

Second, for a given amount of money to be used on redistribution, a UBI likely shifts money away from the very poor. To oversimplify, if you have \$50 billion to alleviate poverty, the targeting approach followed in most countries today might use the \$50 billion to help the poorest/sickest 25 million people and give them the equivalent of \$25,000 of benefits each. With a broad-based UBI, the same \$50 billion would be spread out. It might involve, say, 100 million people getting \$5,000 each. Perhaps a UBI could change the total taste for redistribution in a society—leaving the most disadvantaged people with the same amount and upping the total amount spent—but for the UBI to not end up more regressive than the current system necessarily entails greater amounts of public funds.

Third, the conception of the UBI as a replacement for a myriad of other in-kind transfers and safety net programs forgets the historical origins of that safety net. Fundamentally, the in-kind safety net exists today because rich societies are not comfortable with grievously injured people coming into a hospital but being turned away if they do not have money or letting kids go hungry because their parents cannot afford to feed them, and so on. Converting to a UBI and abolishing the in-kind safety net will lead to a situation where some people will blow their UBI money in unsympathetic ways—gambling, drugs, junk food, Ponzi schemes, whatever. Those people will then

come to emergency rooms or their kids will be hungry and by the rules, they will be out of luck. That is what their UBI income was supposed to cover. But the fact that advanced economies evolved an in-kind safety net in order to avoid this situation makes me think that enforcing “UBI discipline” and replacing the safety net with a straight transfer would require rather extraordinary changes in the psyche of people in the advanced economies.

11.4 Policy Responses to AI beyond Jobs: Pricing, Data Property Rights, and Antitrust

Just as the impact of AI goes far beyond just the impact on employment, the policy response to AI raises all sorts of other considerations, as well.

One is the perennial back-and-forth over the power of buyers versus the power of sellers in pricing. The same issue arose with the initial rise of e-commerce—the new online data on customers allowed new forms of price discrimination and market power but the ease of comparison shopping reduced search costs and promoted competition (e.g., Brown and Goolsbee 2002). So far, the power of the AI technology seems overwhelmingly to have been used by sellers. If they can individualize market and price discriminate with it, margins will likely rise. But consumers will likely push back. They may find technological solutions to use AI to thwart merchants. But a more straightforward response might be to follow past practice and start making various behaviors and practices illegal. This could include restrictions on consumer privacy and the ways that companies can use customer information. It might manifest as an argument over property rights in the sense of who owns the consumers’ data and what level of consent it requires to use it, or might involve rules against various types of price discrimination. Regardless of the form, these issues of pricing and data seem like they will be a central area of policy in an AI-centric world.

The second thing about an AI economy is that the fixed-cost/economies of scale seem pretty significant, and in many cases there are also often network externalities and switching costs on the demand side of these industries. All of these seem to portend the possibility of many industries having a winner-take-all market structure or the continued rise of “platform” competition rather than conventional competition. If so, the rise of AI is likely to usher in a renewed emphasis on antitrust policy in much the same way the original Gilded Age consolidation of industry did before.

11.5 Conclusion: Will Robots Take Over Policy, Too?

The organizers of the volume also asked us to consider whether AI will enhance or even replace the jobs of policy makers—whether improvements in machine learning and AI could be used on the policy-making process itself. Personally, I do not think so because the most important policy mat-

ters are at their heart not issues of prediction. The technology may improve our ability to predict responses, but it does not help us balance interests or engage in politics. We already know, for example, a great deal about the fiscal implications for social security of the aging population. Artificial intelligence might improve our ability to predict revenue outcomes of various policy options, say. That has not been the problem with addressing social security. It has always been about choosing between options and making value judgments. The kinds of problems that AI helps with are those where large amounts of past data to inform the decision. Conditions with small samples or where the conditions are very different than in the past will be much less machine learnable. For small bore issues, AI may improve policy accuracy—what conditions should cause regulators to raise their estimated probability that a bank’s loans will start to default, for example. For bigger issues, though, like whether the Federal Reserve should raise interest rates or whether we should cut high-income people’s taxes—I have my doubts about what AI can contribute.

It is also sure to increase the attention paid to business practices of large AI platforms—their pricing, their use of personal data on customers, their behavior toward competitors, and the continuing consolidation of market power. Each of these is likely to become a major policy battleground of the future. For the time being, though, the job of policy makers themselves seem relatively safe . . . for now.

References

- Autor, David. 2015. “Why Are There Still So Many Jobs? The History and Future of Workplace Automation.” *Journal of Economic Perspectives* 29 (3): 3–30.
- Autor, David, and A. M. Salomons. 2018. “Is Automation Labor-Displacing? Productivity Growth, Employment, and the Labor Share.” *Brookings Papers on Economic Activity* 2018, Spring. <https://www.brookings.edu/bpea-articles/is-automation-labor-displacing-productivity-growth-employment-and-the-labor-share/>.
- Bresnahan, Timothy, and Robert Gordon, eds. 1997. *The Economics of New Goods*. Chicago: University of Chicago Press.
- Brown, Jeffrey, and Austan D. Goolsbee. 2002. “Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry.” *Journal of Political Economy* 110 (3): 481–507.
- Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton.
- Goolsbee, Austan D., and Peter J. Klenow. 2006. “Valuing Consumer Goods by the Time Spent Using Them: An Application to the Internet.” *American Economic Review, Papers and Proceedings* 96 (2): 108–13.
- . 2018. “Internet Rising, Prices Falling: Measuring Inflation in a World of E-Commerce.” *American Economic Review, Papers and Proceedings* 108 (5): 488–92.

- Gordon, Robert. 2016. *The Rise and Fall of American Growth: The US Standard of Living since the Civil War*. Princeton, NJ: Princeton University Press.
- McKinsey Global Institute. 2017. *Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation*. December. McKinsey & Co. Accessed Apr. 26, 2018. <https://www.mckinsey.com/~/media/McKinsey/Global%20Themes/Future%20of%20Organizations/What%20the%20future%20of%20work%20will%20mean%20for%20jobs%20skills%20and%20wages/MGI-Jobs-Lost-Jobs-Gained-Report-December-6-2017.ashx>.
- Jaravel, Xavier. 2017. "The Unequal Gains from Product Innovations: Evidence from the US Retail Sector." Unpublished manuscript, London School of Economics. April.
- Mokyr, Joel. 2014. "Secular Stagnation? Not in Your Life." In *Secular Stagnation: Facts Causes and Cures*, edited by Coen Teulings and Richard Baldwin, 83–89. London: CEPR Press.
- Varian, Hal. 2013. "The Value of the Internet, Now and in the Future." *Economist*, Mar. 10. Accessed Apr. 26, 2018. <https://www.economist.com/blogs/freexchange/2013/03/technology-1>.