PUBLIC POLICY IN AN AI ECONOMY

Austan Goolsbee

University of Chicago Booth School of Business

December 20, 2018

I wish to thank the participants at the NBER Artificial Intelligence conference for helpful comments.

This conference has brought together a mix of technology and economics people to think broadly about the role of Artificial Intelligence in the economy and this short paper 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 ended up focused on jobs and robots. Ruminations by public figures like Bill Gates, Stephen Hawking and Elon Musk have stoked fears that robots will destroy our jobs (and then 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 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 (insert literature cites *** Brynjolfsson and Mcafee, Autor ***, Mokyr **). 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. A major survey across industries by McKinsey Global Institute (McKinsey, 2017) argues that 73 million jobs could 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 in many areas. If the AI helps us diagnose medical problems better, improve our highway safety, give 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 (insert literature cites: see Brynjolfsson and Smith, ***; Varian, ***; Goolsbee and Klenow, ***).

That's 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 we could get AI to change that equation.

This paper 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.

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%. Over time, people adjust. They move. They get skills. The long-run impact of labor saving technologies have 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 before. That really isn't different. If the fear is that AI fundamentally is different because it will begin to replace the types of jobs that have never been automated before like higher-skill, higher-education jobs, the 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. If it all happens at once is when the adjustment problems are worse. Take the much discussed case of autonomous cars. There were about 3.5 million truck, bus, and taxi drivers in 2015. Say that every one of them were lost due to advances in self-driving car technology.

If this loss takes place over 15 years, this would average a little over 19,000 per month. Compare that to the fact that in 2017 the 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 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 much harder to adjust to in the immediate term.

So it's worth considering what influences the speed of adoption. A key determinant will be how good the AI is. But many analysts seem to view that is the only thing that will determine adoption rates. It's worth considering two other factors, at least—prices and adjustment costs.

First, many of these AI innovations are expensive capital outlays up front and that alone may slow their adoption for some time. Ride-share drivers, for example, by various 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 up front fixed costs assuming future prices will rise to cover them?

Second, 'better' doesn't always mean faster. Economists have shown automated stock picking through index funds superior to active management enough times to be a cliché yet people still hold trillions in inefficient, high-fee funds. Millions of people have mortgages with higher than market interest rates 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 driving for themselves.

Third, in an important sense, we know that AI is only 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 human. They started it in the U.S.

and had to shut it down on the first day because it became so abusive and offensive. All it did was mirror 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. 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 doesn't 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 (see Jaravel, 2016 for more discussion about unequal gains from product innovations).

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

CHALLENGES FOR A UNIVERSAL BASIC INCOME RESPONSE TO JOB MARKET DISPLACEMENT

So now suppose that the arguments above prove wrong and 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 for more. In the original 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 US.

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, at least in most cases, 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 is an advantage. In a world like the one we are in now, offering a basic income will likely cause a sizable drop in labor market participation by low wage groups. To the extent that non-participation in exactly that segment of the labor force is already viewed as a problem, the UBI would likely make things worse.

Second, for a given amount of money to be used on redistribution, UBI likely shifts money away from the very poor. In other words, if you have $50B to alleviate poverty, the targeting approach followed in most countries today might use the $50B to help the poorest/sickest 25m people and give them the equivalent of $25,000 of benefits each. With a broad-based UBI, the same $50B would be

spread over many more people. It might involve 100m people getting $5000 each, say. 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 seems not to pay much mind to the origins of that safety net. Fundamentally, the in-kind safety net exists today because rich societies are fundamentally not comfortable with letting people come to the hospital and be turned away to die because they don't have money or letting kids go hungry because their parents do not have the money to feed them, and so on. Converting things to a UBI and getting rid of the in-kind safety net will lead to a situation in which, even if among a small share of UBI recipients, SOME people will blow their money in unsympathetic ways—gambling, drugs, junk food, Ponzi schemes, whatever. And now those people will come to the emergency room or their kids will be hungry and by the rules, they will be out of luck. That's what they were supposed to have used their UBI for. That society has an in-kind safety net now in order to avoid this circumstance makes me think that enforcing "UBI-discipline" and replacing the safety net with a straight transfer would require extraordinary changes in how society functions.

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 ecommerce. The new data on customers allowed new forms of price discrimination and market power. The ease of comparison shopping reduced search costs and reduced market power. 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 and consumers will likely push back. They might find technological solutions but the most common response in the past is toughening consumer protection laws. 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 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 anti-trust policy in much the same way the original gilded age consolidation of industry did before.

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 matters are at their heart not issues of prediction. The technology may improve our ability to predict responses but not 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 hasn't been the problem with addressing social security. It has been about choosing between tradeoffs 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 very small samples or where the conditions are very different than in the past will be much less machine learnable. For small bore issues—should regulators be nervous about a loan portfolio with X characteristics?—AI may improve policy accuracy. For bigger issues—should the Federal Reserve raise interest rates in December?—I have my doubts.

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 policymakers themselves seem relatively safe…for now.

Bibliography

Autor, David

Brown and Goolsbee

Brynjolfsson, Erik and Andrew McAfee

Brynjolfsson, Erik and Michael Smith

Goolsbee and Klenow

Gordon, Robert

Mckinsey Global Institute

Mokyr, Joel

Varian, Hal