

NBER WORKING PAPER SERIES

ECONOMIC POLICY FOR ARTIFICIAL INTELLIGENCE

Ajay K. Agrawal
Joshua S. Gans
Avi Goldfarb

Working Paper 24690
<http://www.nber.org/papers/w24690>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2018

Thanks to participants at the 2017 NBER Conference on the Economics of Artificial Intelligence and at the 2018 NBER Innovation Policy and the Economy Conference for ideas and comments. We thank the Sloan Foundation for financial support of the NBER Economics of Artificial Intelligence initiative. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w24690.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Ajay K. Agrawal, Joshua S. Gans, and Avi Goldfarb. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Economic Policy for Artificial Intelligence
Ajay K. Agrawal, Joshua S. Gans, and Avi Goldfarb
NBER Working Paper No. 24690
June 2018
JEL No. L86,O3

ABSTRACT

Recent progress in artificial intelligence (AI) – a general purpose technology affecting many industries - has been focused on advances in machine learning, which we recast as a quality-adjusted drop in the price of prediction. How will this sharp drop in price impact society? Policy will influence the impact on two key dimensions: diffusion and consequences. First, in addition to subsidies and IP policy that will influence the diffusion of AI in ways similar to their effect on other technologies, three policy categories - privacy, trade, and liability - may be uniquely salient in their influence on the diffusion patterns of AI. Second, labor and antitrust policies will influence the consequences of AI in terms of employment, inequality, and competition.

Ajay K. Agrawal
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
CANADA
and NBER
ajay.agrawal@rotman.utoronto.ca

Avi Goldfarb
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
CANADA
and NBER
agoldfarb@rotman.utoronto.ca

Joshua S. Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto ON M5S 3E6
CANADA
and NBER
joshua.gans@gmail.com

1. Introduction

Artificial intelligence (AI) technologies have advanced rapidly over the last several years. As the technology continues to improve, it may have a substantial impact on the economy with respect to productivity, growth, inequality, market power, innovation, and employment. In 2016, the White House published four reports emphasizing this potential impact.² In 2018, it announced a set of policies that signals a ‘hands off’ approach to AI.³ The governments of China and France have made it a priority. Other governments around the world have dedicated substantial resources to investing in the technology and preparing for its impact.

Pessimistic views of the impact of AI on society are widespread. Public figures including Elon Musk and Stephen Hawking have warned that AI could lead to a handful of companies dominating society, few jobs left for humans, and increasing inequality. Vladimir Putin ominously predicted, “[w]hoever becomes the leader in this sphere will become the ruler of the world” (Vincent 2017). Among economists, Robert Gordon (2016) has a different kind of pessimistic view. Rather than worrying about AI leading to rapid change and upending society, he argues that the inventions of the future are unlikely to match the impact of the inventions of the period from 1870 to 1970. More optimistically, AI could enhance productivity so dramatically that people have plenty of income and little unpleasant work to do (Stevenson 2018).

Under either the optimistic or the pessimistic view, policy will shape how AI impacts society. In

² The reports are: (1) Jason Furman, “Is This Time Different? The Opportunities and Challenges of Artificial Intelligence” (remarks at AI Now, New York University, July 7, 2016), https://obamawhitehouse.archives.gov/sites/default/files/page/files/20160707_cea_ai_furman.pdf; (2) Executive Office of the President, “Artificial Intelligence, Automation, and the Economy,” December 2016, <https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF>; (3) Executive Office of the President, National Science and Technology Council, and Committee on Technology, “Preparing for the Future of Artificial Intelligence,” October 2016, https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf; (4) National Science and Technology Council and Networking and Information Technology Research and Development Subcommittee, “The National Artificial Intelligence Research and Development Strategic Plan,” October 2016, https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/national_ai_rd_strategic_plan.pdf.

³ <https://www.scienceaf.com/white-house-trump-administration-ai-committee-keep-america-first>

this chapter, we discuss a variety of government policies that are likely to influence the diffusion and consequences of AI. In doing so, we draw heavily on the ideas presented at the NBER Conference on Artificial Intelligence (held in September 2017 in Toronto) and the associated conference volume (Agrawal, Gans, and Goldfarb 2018a). We first define AI and argue that it is likely to be the next general purpose technology (GPT) characterized by pervasive use across a wide variety of sectors and technical dynamism (Bresnahan and Trajtenberg 1995). We next discuss the policies that are likely to influence the speed of diffusion of AI and the form that the technology takes, with an emphasis on industry regulation. We then discuss policies related to addressing the potential consequences of the diffusion of AI, emphasizing productivity and labor market changes. In the final part of this section, we discuss the potential for AI to generate new antitrust concerns. We conclude with a summary of the key policy issues.

Artificial Intelligence as a General Purpose Technology

The economic impact of computer intelligence that can do what humans do would be extraordinary. At the NBER Conference on AI in September 2017 Kahneman (2018) noted, “I do not think that there is very much that we can do that computers will not eventually be programmed to do.” That said, the technology that has generated the recent excitement around AI is not true artificial general intelligence. Instead, the majority of recent advances have come from a field of AI called “machine learning,” which is a branch of computational statistics.

The Oxford English Dictionary defines artificial intelligence as “the theory and development of computer systems able to perform tasks normally requiring human intelligence.” This definition is both broad and fluid. There is an old joke among computer scientists that artificial intelligence defines what machines cannot yet do. Before a machine could beat a human expert at chess, such a win would mean artificial intelligence. After the famed match between IBM’s Deep Blue and Gary Kasparov, playing chess was called computer science and other challenges became artificial intelligence.

Computer chess and other early attempts at machine intelligence were primarily rules-based, symbolic logic. It involved human experts generating instructions codified as algorithms (Domingos 2015). By the 1980s, it became clear that outside of very controlled environments, such rules-based systems failed. They did not cope with the complexity of most applications and AI research funding largely dried up (Markov 2015).

More recently, a different approach has taken off: Machine learning. The idea is to have computers “learn” from example data. Humans conduct many tasks that are difficult to codify. For example, humans are good at recognizing familiar faces, but we would struggle to explain this skill. By connecting data on names to image data on faces, machine learning solves this problem by predicting which image data patterns are associated with which names.

Economists looking at a machine learning textbook will find many familiar topics, including multiple regression, principal components analysis, and maximum likelihood estimation, along with some that are less familiar such as hidden Markov models, neural networks, deep learning, and reinforcement learning. Even these are not entirely new. Igami (2018) highlights similarities between these topics and structural econometrics models such as Rust (1987). Unlike the typical emphasis of econometricians on causal inference, machine learning focuses on prediction. As machine learning pioneer Geoffrey Hinton put it, “Take any old problem where you have to predict something and you have a lot of data, and deep learning is probably going to make it work better than existing techniques”.⁴ In our own work (Agrawal, Gans, and Goldfarb 2018b,c,d), we explore the consequences of a drop in the quality-adjusted cost of prediction on organizations and on the nature of work. We define prediction as the ability to take known information to generate new information. Better prediction is likely to have widespread consequences because prediction is fundamental to decision-making.

⁴ <https://www.youtube.com/watch?v=2HMPRXstSvQ>. Accessed May 22, 2018.

Therefore, like the steam engine, electrification, and the internet, AI—defined by machine learning—is likely to be a general purpose technology (GPT). As Brynjolfsson, Rock, and Syverson (2018) argue, machine learning is likely to impact a wide range of sectors, and there is no doubt that technological progress has been rapid. They note that error rates in image recognition improved by an order of magnitude between 2010 and 2016. There have been rapid improvements in areas from language translation to medical diagnosis. As a GPT, it is likely to be an enabling technology that opens up new opportunities. For example, while electric motors did reduce energy costs, the productivity impact was largely driven by increased flexibility in the design and location of factories (David 1990). This process of co-invention is costly and slow (Bresnahan and Greenstein 1996).

As with other GPTs, this suggests that the impact of AI on productivity and living standards will take time to arrive. In responding to skeptics such as Gordon (2016), Brynjolfsson, Rock, and Syverson (2018) argue that the impact of information technology is still likely in an early phase. Their argument draws on examples from the past. Figure 1 compares labor productivity growth over the first fifty years of portable power and the first 50 years of widespread application of information technology. The figure suggests that the labor productivity effects of portable power may have taken close to 50 years, and so some patience for waiting for a measurable impact of information technology generally, and AI in particular, is warranted.

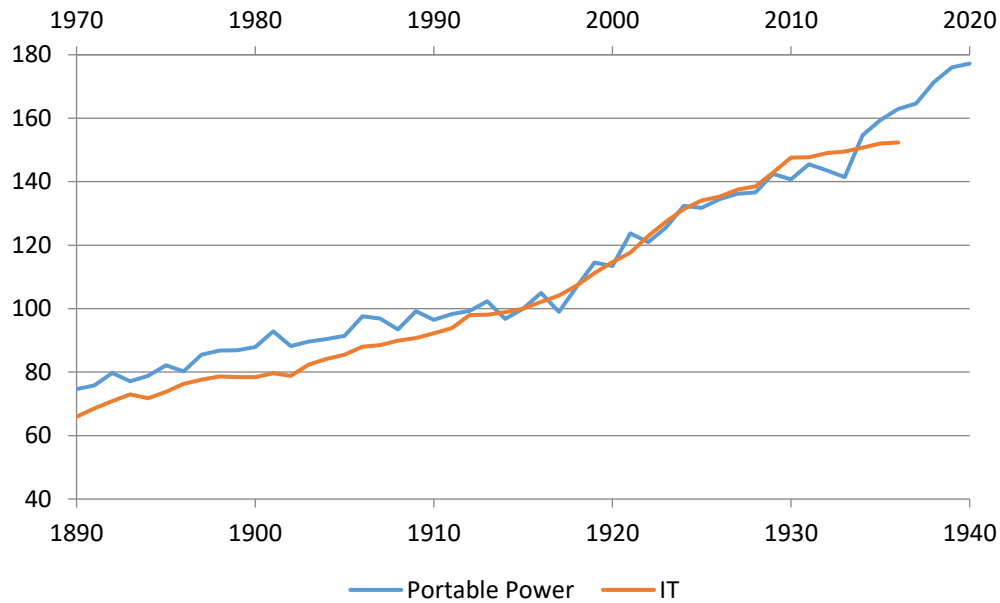


Figure 1: Labor Productivity Growth in the Portable Power and IT Eras.
 Source: Brynjolfsson, Rock, and Syverson (2018). Figure 8.

Many of these new opportunities will be in science and innovation. Cockburn, Henderson, and Stern (2018) empirically demonstrate the widespread application of machine learning in general, and deep learning in particular, in scientific fields outside of computer science. For example, Figure 2 shows the publication trend over time for three different AI fields: Machine learning, robotics, and symbolic logic. For each field, the graph separates publications in computer science from publications in application fields. The dominant feature of this graph is the sharp increase in publications that use machine learning in scientific fields outside computer science. Along with other data presented in the paper, they view this as evidence that AI is a GPT in the method of invention. It will, therefore, have a widespread impact on the economy, accelerating growth.




Figure 2: Publications in Computer Science versus Application Journals, by AI Field
 Source: Cockburn, Henderson, and Stern (2018). Figure 4

Aghion, Jones, and Jones (2018) demonstrate that if AI is an input into the production of ideas, then it could generate exponential growth even without an increase in the number of humans generating ideas. Put differently, recent empirical work suggests that scientific ideas may be getting harder to find (Bloom et al 2017). If AI is an input into the production of ideas, this would provide a way out of this potential driver of slowing productivity growth. Instead, growth could accelerate.

Before concluding this section on AI as a GPT, it is important to note that the above discussion emphasizes advances in machine learning, leaving the impact of the still-speculative technology of artificial general intelligence aside. While Kahneman noted the possibility of the technology, it is an open question whether these advances in machine learning are sufficient to get us to an artificial general intelligence in the near term. Subject matter experts disagree. Ben Goertzel, president of the Artificial General Intelligence Society expects just seven years. Reinforcement learning pioneer Richard Sutton provided a median guess that artificial general intelligence will arrive by 2040. Gary Marcus, often seen

as an AI skeptic, provided a confidence interval of 30 to 70 years.⁵ Many others say never. If artificial general intelligence is imminent, then a key policy challenge will be for the institutions to keep up with the abilities and needs of a world with machine intelligence. Such rapid change would suggest that relying on the historical experience of past GPTs is unlikely to be useful guide to the impact of AI. In contrast, if machine learning will need to be combined with many other advances in order to be applicable across a wide range of applications, then the above discussion that emphasizes the GPT model, the importance of co-invention, and the use of history as a guide is appropriate. Given the still-speculative nature of artificial general intelligence, in the discussion that follows we focus on the impact of machine learning and emphasize its potential as a GPT.

2. Policies that will influence the diffusion of AI

Policy will have an impact on the rate of improvement in AI, the speed of diffusion, and the nature of the technology. Most obviously, as with other technologies, policies that provide research support for AI are likely to accelerate technological progress as are policies that strike the appropriate balance with respect to intellectual property (Williams 2016; Scott Morton and Shapiro 2016).⁶ However, three other policy categories - privacy, trade, and liability - are likely to impact the diffusion of AI in ways that are different than other technological innovations. We discuss each in turn.

⁵ Market for Intelligence Conference 2017. October 26, Toronto. Video at https://www.youtube.com/watch?v=RxLIQj_BMhk.

⁶ For intellectual property the main trade-off is between creating incentives to innovate by protecting intellectual property against allowing competition to reduce prices and increase consumer surplus. There is also a trade-off between incentives for the inventor to innovate and ability to create follow-on innovation (Green and Scotchmer 1995).

Privacy

A key input into AI is data. Machine learning uses data to make predictions about what individuals may desire, be influenced by, or do. This creates privacy concerns. Tucker (2018) emphasizes that privacy is challenging for three reasons: (1) *cheap storage* means that data may persist longer than the person who generated the data intended, (2) *non-rivalry* means that data may be repurposed for uses other than originally intended, and (3) *externalities* mean that data created by one individual may contain information about others. Thus, the collection and usage of data to make predictions has the potential to harm individuals if they are not fully aware of how their data is being collected and used. As data collection and analysis has improved, policy-makers have turned increasing attention to regulating the collection, storage, and usage of data. In data-driven industries, privacy regulation impacts the rate and direction of innovation (Goldfarb and Tucker 2012). Too little privacy protection means that consumers may be unwilling to participate in market transactions where their data are vulnerable. Too much privacy regulation means that firms cannot use data to innovate.

The existing empirical evidence on the effect of privacy regulation on innovation does not focus on AI specifically. Broadly, however, the evidence to date indicates that most attempts at government-mandated privacy regulation lead to slower technology adoption and less innovation (Goldfarb and Tucker 2011, Miller and Tucker 2011, Kim and Wagman 2015). For example, Goldfarb and Tucker (2011) provide evidence that suggests European regulations around the use of online tracking technologies reduced the effectiveness of European banner advertising by 65% relative to US banner advertising. The results suggest that this is driven by a reduced ability of European advertisers to use data to target advertising. To the extent that many of the most valuable companies today are advertising platforms (including Google, Facebook, and Baidu), the benefit of privacy may have come at a cost of local companies in these industries. In healthcare, Miller and Tucker (2009) demonstrate how privacy restrictions slowed the diffusion of electronic medical records. Miller and Tucker (2011) show that electronic medical records are

a life-saving technology, providing substantial benefits to health. Combined, the Miller and Tucker papers suggest that privacy regulation in healthcare can reduce health outcomes. Together, these papers show that privacy regulations slow the diffusion of technology that relies on data.⁷ In some sense, that is to be expected. If privacy is a fundamental human right, then we should be willing to trade off something in exchange for it. Otherwise, there is little value in considering it a right. While research has not yet studied how privacy policy affects AI innovation in particular, it is clear that privacy policy will have an influence on the diffusion of AI.

Trade

Many countries view investments in AI as strategic. Chinese investments in AI led to a marked increase in the share of research papers at major conferences written by Chinese scientists, while the share of papers from European, Australian, and Canadian scientists fell substantially (though not the absolute numbers because the total number of papers increased). AI-related trade policy refers to behind-the-border policies that are increasingly included in trade agreements. For example, privacy policy varies across countries. Therefore, privacy policy is likely to affect the relative diffusion of AI across countries. If the optimal privacy policy for the diffusion of AI is relatively lax, then countries with lax privacy policies may develop a short-term advantage. Economies of scale in data and the potential for learning-by-doing on the part of researchers create incentives for a “race to the bottom” in privacy policy, in which countries reduce privacy in order to get ahead of each other in AI. Prior research (e.g., Davies and Valdlamannai 2013, Fredriksson and Milliment 2002), identified such races in labor and environmental policies and the same forces may operate in privacy policy. To the extent that Europe continues to have stricter privacy

⁷ In contrast, voluntary privacy-protecting actions by firms do not seem to have the same kinds of effects and could be either neutral or innovation-enhancing depending on the context (Goldfarb and Tucker 2012). For example, relative to its competitors, Apple has invested in privacy and there is no evidence to suggest that this has been harmful (Agrawal, Gans, and Goldfarb 2018d).

policies than the US, and the US stricter than China, this could create a disadvantage for European companies relative to US relative to China for accessing the data needed to build many types of AI.

In terms of trade policy, trade agreements could mitigate races to the bottom by specifying international privacy standards, just as trade agreements sometimes specify labor and environmental standards. Such privacy standards would still affect the rate and direction of innovation, but in a way that would not give any particular country an advantage. It is not clear if or how this would impact overall economic growth and productivity. Those countries with stricter privacy protections may focus on other industries. In the context of patents, for example, Moser (2005) found that the weaker patent protection of some countries in the 19th century such as Switzerland meant that innovators in those countries focused on innovations which could be protected by other means, such as trade secrets: Swiss clock making skills could not be easily imitated.

Goldfarb and Trefler (2018) argue that strategic trade models show that races to the bottom are only relevant if there are likely to be country-level rents from the technology. Other behind-the-border policies that are relevant to trade and have been discussed in the context of trade agreements such as TPP-11, include data localization rules, limited access to government data, industry regulations such as those around the use of drones, and forced access to source code. The key question when evaluating any of these policies is whether there are local spillovers from AI, or whether the benefits of the technology diffuse sufficiently widely that free entry mitigates any potential rents. More generally, AI is only relevant to effective trade policy to the extent that there are country-level rents to be gained.

The implications of some government actions imply that some anticipated AI-related rents will be large. Investments by certain governments, such as China and France, suggest that these countries expect large localized rents. Many of the most valuable private companies in the world have made large investments in AI. Machine learning improves with more data, and so larger companies and larger countries could get the most benefit. At the same time, the rapid geographic diffusion of artificial

intelligence research and innovation suggests that rents might not be local. Commercial applications have often occurred far from the location of the inventions. Thousands of people are trained every year to apply machine learning algorithms using open source code. Furthermore, in a technical sense, there are decreasing returns to scale in data because prediction accuracy improves in the square root of N (Bajari, Chernozhukov, Hortacsu, and Suzuki 2018).⁸ Overall, the industrial application of AI is still in its infancy, and so the extent to which AI-related rents will be localized is unclear. Nevertheless, given the investments governments around the world are making in AI, a race to the bottom in privacy is possible, and it is a likely issue in future trade agreements.

Liability

Galasso and Luo (2018) emphasize that rapid advancements in the field of artificial intelligence and robotics have led to lively debates over the application of tort law to these technologies. A tort is an action that causes harm or loss, resulting in legal liability for the person who commits the act. The purpose of the tort system is to deter people and companies from injuring others, and to compensate injured parties. Porter (1990) argues that liability risks deter innovation in the United States. In particular, because the outcome of tort cases can be uncertain, and because payouts can be high, he argues that firms may avoid innovations that—though likely to be beneficial—carry some tort risk. The empirical evidence on liability is more mixed. Viscusi and Moore (1993) document that potential liability is positively correlated with product R&D at the industry level (except in cases of very high risk), likely because it increases investment in safety-improving technology. Galasso and Luo (2017) find that reduced tort risk reduces medical device innovation.

⁸ Importantly, this last point assumes that the return to scale on the business side match the returns to scale of the technology. If being slightly better than the competition leads to dominance in market share—there can still be increasing returns to data in terms of the economic opportunity.

For AI, getting the balance right on liability will be important to its diffusion (Galasso and Luo 2018). In the absence of clear liability rules for AI products that involve many producers, firms may not invest in the technology. Consider autonomous vehicles. Autonomous vehicle production and use will likely depend on many different firms, including sensor manufacturers, telecommunications providers, vehicle manufacturers, and AI software developers. Without the clear assignment of who is liable for faulty car performance, all may hesitate to invest in the technology. Furthermore, with less room for consumers to take precautions, the liability will shift to producers. Similarly, in health care, robotic surgery and machine learning-based diagnosis may shift liability from health care providers to device manufacturers. In order to address these risks, the European Parliament adopted a resolution in February 2017 for EU-wide legislation to regulate “sophisticated robots, bots, androids and other manifestations of artificial intelligence” (European Parliament, 2017). The US House passed liability legislation for autonomous vehicles in September 2017.

Algorithmic bias is another way that liability could affect the nature and diffusion of AI. While much evidence suggests that AI predictions are less biased than human predictions (Hoffman, Kahn, and Li 2018), AIs are still trained on data generated by humans and so predictions based on AIs will perpetuate those biases. Furthermore, it is easier to audit AI-based decisions than human decisions and this opens up the possibility of an increase in liability claims.

Lambrecht and Tucker (2018) show that algorithms can be biased, even when firms do not intend it and the training data is not obviously biased. They showed that advertisements for “STEM” education on Facebook and twitter were shown more to men than women. Given that women are under-represented in STEM fields, the outcome was biased. Lambrecht and Tucker showed that this difference was driven by the price of advertising and not any explicit bias. In particular, dynamic pricing algorithms priced advertising to women higher than advertising to men and so, in response, the customer-facing advertising algorithm, programmed to maximize clicks per dollar, saved the advertiser money by showing

more ads to men. Ben Edelman (as noted in Agrawal, Gans, and Goldfarb 2018d) emphasized that firms face liability risk even in such situations when the bias is unintended. Such risk could serve as a deterrent to the adoption of AI in situations where algorithms might be biased, even if human decision-makers would be equally or more biased because the algorithms may be more easily audited in the courts.

3. Policies that address the consequences of the diffusion of AI

In the previous section, we emphasize how policy may enhance the rate of diffusion of AI. However, there is not universal agreement that diffusion is a good idea. As noted earlier, there are both pessimistic and optimistic views of the consequences of AI. There are two distinct elements of the pessimistic view. First, AI will replace jobs leaving little for humans to do. Second, AI will not be as influential as the technologies that diffused between 1870 and 1970 and so it will have relatively little impact on society. While both issues share a pessimistic vision of the future, Mokyr (2017) notes that “The good news is that those pessimistic predictions cannot both be right. The even better news is that they can both be wrong.” If AI is so efficient that it replaces humans, then it will be transformative. If it is not transformative, then it will have little impact by definition.

Nevertheless, our view of the economic discussion around AI is that, to the extent that AI is likely to be a productivity-enhancing technology, the worry is that it will diffuse too slowly not too quickly (Furman 2018, Goolsbee 2018). At the same time, the impact will not affect all people and all firms equally. Stevenson (2018) summarizes the key issues as “two separate questions: there’s an employment question, in which the fundamental question is can we find fulfilling ways to spend our time if robots take our jobs? And there’s an income question, can we find a stable and fair distribution of income?” We discuss the employment question and then the income distribution question.

AI and Jobs

A common worry about the diffusion of new technologies is the impact on jobs. If a technology does what a worker does, but better, then what will that worker do? Acemoglu and Restrepo (2018) emphasize that a job can be decomposed into tasks, and so the impact of AI and automation on jobs depends on the tasks the AI will replace. They emphasize a displacement effect, as machines take over human tasks and an expansion effect as growth leads to the creation of new tasks in which labor has a comparative advantage over machines. Like others (Furman 2018, Goolsbee 2018, Stevenson 2018, Bessen 2018, etc.), the paper suggests an optimistic long run message; however, a key theme is that adjustment may be slow and so there is likely to be a short and medium-term mismatch of skills and technologies.

Thus, the key policy questions with respect to AI and jobs relate to the business cycle and education policy. For coping with the business cycle, AI-related policy is not likely to be notably distinct. The potential for layoffs concentrated in location and time is not unique to AI. It was a feature of factory automation and the mechanization of farming and may also apply to other emerging technologies such as 3D printing.

For education policy, Trajtenberg (2018) highlights three types of skills that are likely to be needed as AI diffuses: analytical and creative thinking, interpersonal communication, and emotional control. Agrawal, Gans, and Goldfarb (2018b,d) emphasize the role of human judgment, defined as the ability to identify what to do with a prediction (e.g., determine the payoff function in the context of decision making). Judgment is the skill of knowing the objectives of an organization and translating that into data that can be collected. Similarly, Francois (2018) emphasizes the skill of telling the machines what to optimize. This requires understanding both the capabilities of the machine and the goals of the organization. It represents a combination of technical skills and social science. An alternative possibility is social skills, as discussed by Deming (2017). If machines are increasingly going to do prediction and

technical tasks, what might be left for humans involves engaging, motivating, and comforting other humans.

Whether good jobs and continued economic growth will require technical skills, social skills, or something else depends on the sectors that make up the largest share of the economy. To the extent AI increases productivity, Baumol's (1967) work on "cost disease" provides insight into the limits of growth: Economic growth is constrained by important things that are hard to improve and so sectors that see rapid productivity growth will see their share of GDP decline. In other words, as emphasized by Aghion, Jones, and Jones (2018), the sectors that adopt AI most quickly and effectively may play a decreasing role in the economy over time. As noted above, this raises the potential for short run job losses, and the question of what skills will be most needed for strong economic growth. The fastest growing sectors in terms of employment may not be the sectors that use AI best, and so many jobs of the future could be in the sectors where AI has the least impact.

It is not clear exactly how this will play out. It might mean that Deming's emphasis on social skills proves to be most important in the future; however, Kahneman (2018) argues that machines are likely to be more emotionally intelligent than humans because it is relatively easy to predict human reactions to certain stimuli and because machines won't get grumpy, hungry, or emotional. Therefore, it might mean that other skills are key.

A different issue in education policy is whether the current model in which we front load education early in life will remain optimal. If AI enables machines to "learn" such that humans are faced with continual change in the skills required to be employable, then we may need to rethink the structure of education (continual). Education policy will need to adapt to AI both in terms of the skills taught and the structure of delivery. It might suggest a need to teach skills related to telling an AI what to predict. It might suggest a model of continual learning through adulthood. Without further research, the nature of the adaptation remains an open question.

In summary, the near consensus among economists is optimism for the long run prospects of jobs and growth; however, the long run can be a very long time. There are short-run considerations with respect to the business cycle and a mismatch between skills and technology. Education policy in particular is likely to be important for making the most of the diffusion of AI.

Stevenson (2018) also highlights a broader policy question around the meaning of work. Over the past century, many societies choose to spend their increased wealth on leisure. Child labor was reduced, the hours worked per week were reduced, and retirement became an expected stage of life. If AI increases wealth further, it may open the possibility of further reductions in work as an increase in leisure that is distinct from a loss of jobs. Stevenson emphasizes that a key challenge for society might be to identify meaning in the absence of full-time employment for many people. If we choose to have more leisure time, and we understand it as leisure time, that might be a positive result of AI's diffusion even if the number of hours worked declines.

AI and Inequality

A number of economists have emphasized that AI may increase inequality, even if the long run prospects for jobs and overall growth are positive. There are two broad reasons why AI might lead to an increase in inequality. First, AI—like computers and the internet—might be skill-biased (e.g. Autor, Katz, and Krueger 1998; Akerman, Gaarder, and Mogstad 2015). It is likely to disproportionately increase the wages of highly educated people and might even decrease the wages of the less educated. Furman (2018) emphasized this point, noting that current projections of job losses are much higher for low-wage jobs (see Figure 3) and jobs requiring less education. As Goldin and Katz (2010) emphasized in the context of computing, “[i]ndividuals with more education and higher innate abilities will be more able to grasp new and complicated tools.” To the extent that highly educated people are better at learning new skills, and the skills to succeed with AI will change over time, then the educated will benefit disproportionately. The

most direct policy solution relates to education: If highly educated people benefit most, then the seemingly logical conclusion is to create more highly educated people. However, this logic only works if all people are equally likely to benefit from education. In contrast, if skill-bias is driven by some ability that is correlated with, but not entirely caused by, education, then a policy recommendation for more education might fail. Again, in the context of AI, this is an open question that requires further research.




Figure 3: Probability of automation by Occupation Median Hourly Wage.
Source: Furman (2018). Figure 4

It is possible that the opposite might occur. Recent advances in AI should be seen as improvements in prediction technology. For many of the most highly paid jobs today, such as medical doctor and financial analyst, prediction is a core task. For example, a key role for most doctors is diagnosis. Many of the other aspects of medical care are undertaken by lower-paid medical occupations. Diagnosis is a prediction problem. It takes data on symptoms and fills in the missing information of the cause of those symptoms. If prediction is the highest skilled task in many high wage occupations, then the diffusion of AI could lead to de-skilling and reduced inequality.

While possible, we do not think this is likely. In the past, highly skilled workers have been able to learn new skills as technology automates certain aspects of their jobs while less skilled workers have found this more challenging. Accounting provides a useful example. Accountants used to spend much of their time adding columns of numbers. Computers dramatically reduced the time that accountants spent doing arithmetic but accountants learned a new set of skills that involved leveraging the efficient arithmetic that computers provided (Agrawal, Gans, and Goldfarb 2018b). In contrast, manufacturing workers have had a more difficult time adjusting to the automation of factories. While there are a variety of reasons behind this, a key issue is the challenge of learning new skills for people who have not spent their adult lives focused on learning new skills.

A second reason why AI might lead to an increase in inequality relates to an increased capital share in the economy (Piketty 2013). There is plenty of evidence that the labor share of GDP is falling (Autor et al 2017). If AI is a new, efficient form of capital, then it seems likely that the capital share will rise at the expense of labor (Acemoglu and Restrepo 2018; Sachs 2018).

Policies aimed at dealing with the inequality consequences of AI largely involve changes to the social safety net. One policy that has been widely discussed is the taxation of capital. Bill Gates called for a taxation of robots, though standard models suggest that such a policy would lead to less investment, slower productivity growth, and a poorer society overall. Digging into the standard arguments, Stiglitz and Korinek (2018) provide models for the conditions under which taxation of capital could generate reduced inequality without causing economic stagnation. First, they show that as long as there is a necessary but fixed factor of production (such as materials), taxing that factor can enable redistribution without creating distortions. Second, they show that as long as the supply elasticity of capital is sufficiently low, a combination of intellectual property rights and capital taxation can enable redistribution with minimal distortions.

A second policy that has received a great deal of attention is the universal basic income, which would provide a regular, unconditional cash grant to every individual in society. Furman (2018) and Furman and Seamans (2018) emphasizes that a universal basic income is likely to increase inequality because it would go to all members of society regardless of income, in contrast to the current system in which transfers are aimed at the lower half of the income distribution. In addition to this issue, Goolsbee (2018) notes that a universal basic income might reduce labor market participation among low wage groups. Furthermore, he argues that “rich societies are fundamentally not comfortable with letting people come to the hospital and be turned away to die...or letting kids go hungry” even if people don’t have the money for unsympathetic reasons, such as spending on gambling or drugs. He therefore argues that “replacing the existing safety net with a straight transfer is unrealistic.”

In addition to capital taxation and universal basic income, the existing social safety net could be reinforced in order to reduce inequality whether through increased transfers, healthcare expenditures, or more progressive taxation. The AI context is not unique for these policies and each has well-known economic costs and benefits.

Antitrust and AI

Currently, the leading companies in AI are large in terms of revenue, profits, and especially market capitalization (high multiples on earnings). This has led to an increase in antitrust scrutiny of the leading technology firms from governments (particularly the European Commission) and in the press (e.g. *The Economist’s* January 20-27 2018 cover story “The new titans, and how to tame them”). Much of this antitrust scrutiny focuses on the role of these firms as platforms. Here, we focus on the particular features of AI that are relevant to antitrust: most notably the role of data.

Varian (2018) emphasizes that data is a scarce resource that exhibits decreasing returns to scale in a technical sense: Prediction accuracy theoretically increases in the square root of the number of

observations, suggesting a concave relationship between the amount of data and its value in improving predictions. Under this argument, AI is unlikely to generate antitrust concerns (Bajari et al., 2018).

A more dynamic perspective suggests that there may be economies of scale in data in terms of the business value. In particular, if a slight lead in quality leads to a large lead in the number of consumers, then it could create a circle in which a slight lead in data allows a company to collect more and better data, reinforcing that lead and generating dominance over time. Thus, although there may be decreasing returns to scale in a technical sense, there may be increasing returns to scale in the economic value of data. In such situations, the diffusion of AI may suggest an increase in the importance of antitrust enforcement over time.

Finally, a more speculative antitrust concern with respect to AI has been put forward by Ezrachi, and Stucke (2016). They argue that firms will deploy algorithms programmed to learn to set prices in oligopolistic competition with other firms and their algorithms. It is possible that competing AIs—given the goal of maximizing long-term profits—would learn to tacitly collude. This would be achieved with no overt collusion or communication with the AIs overcoming the technical complexity of colluding in certain environments. We stress, however, that this remains speculative. Current ‘game playing’ AI’s still operate in environments with a limited action space (unlike the broad space that would cover pricing, let alone multi-product pricing). A more likely antitrust application is that antitrust authorities might be able to deploy AI techniques to identify collusive behavior.

4. Conclusions

We highlight two types of policy implications in response to advances in AI technology: policies that affect diffusion patterns and policies that address the consequences of diffusion. The most relevant diffusion-related policy categories are privacy, trade, and liability. Policy design will focus on achieving the desired balance between encouraging diffusion without compromising societal values. As AI diffuses, it

will have consequences for jobs, inequality, and competition. Addressing these consequences will be the role of education policy, the social safety net, and antitrust enforcement.

References

Acemoglu, Daron, and Pascual Restrepo. Forthcoming. The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment. *American Economic Review*.

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

Acemoglu, Daron, and Pascual Restrepo. 2018. Artificial Intelligence, Automation and Work. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Aghion, Philippe, Benjamin Jones, and Charles Jones. 2018. Artificial Intelligence and Economic Growth. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, editors. 2018a. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018b. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business School Press.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018c. Human Judgment and AI Pricing. *American Economic Association Papers and Proceedings*.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018d. Prediction Machines, Judgment, and Complexity. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Athey, Susan. 2018. The impact of machine learning on economics. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2017. The Fall of the Labor Share and the Rise of Superstar Firms. Working paper, MIT.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. Computing Inequality: Have Computers Changed the Labor Market? *Quarterly Journal of Economics*, 113(4), 1169–1213.

Autor, David, and Anna Salomons. 2018. Robocalypse Now: Does Productivity Growth Threaten Employment? Working paper, MIT.

Bajari, Patrick, Victor Chernozhukov, Ali Hortaçsu, and Junichi Suzuki. “The Impact of Big Data on Firm Performance: An Empirical Investigation.” NBER Working Paper.

Baumol, William J. 1967. Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis. *American Economic Review*. 57, 415-426.

- Bloom, Nicholas, Charles Jones, John Van Reenen, and Michael Webb. 2017. Are Ideas Getting Harder to Find? NBER Working Paper # 23782.
- Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence, *The Quarterly Journal of Economics*, 117(1), 339–376.
- Bresnahan, Timothy F., and Shane Greenstein. 1996. Technical progress in computing and in the uses of computers *Brookings Papers on Economic Activity: Microeconomics*, 1-78.
- Bresnahan, Timothy F., and M. Trajtenberg. 1995. General purpose technologies ‘Engines of growth’? *Journal of Econometrics* 65, 83-108.
- Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2018). Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Chevalier, Judy. 2018. Antitrust and Artificial Intelligence: Discussion of Varian. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Cockburn, Iain, Rebecca Henderson, and Scott Stern. 2018. The Impact of Artificial Intelligence on Innovation. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Crafts, Nicholas. 2004. Steam as a general purpose technology: A growth accounting perspective, *Economic Journal*, **114**(495), 338-351.
- David, Paul A. 1990. The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *American Economic Review Papers and Proceedings* 80(2), 355-361.
- Ezrachi, Ariel, and Maurice E. Stucke (2016), *Virtual Competition*, Oxford University Press.
- Furman, Jason. 2018. Should We Be Reassured If Automation in the Future Looks Like Automation in the Past? In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence*. University of Chicago Press.
- Furman, Jason, and Robert Seamans. 2018. AI and the Economy. In Josh Lerner and Scott Stern Eds. *Innovation Policy and the Economy* volume 20. University of Chicago Press. Forthcoming.
- Galasso, Alberto, and Hong Luo. 2018. Punishing Robots: Issues in the Economics of Tort Liability and Innovation in Artificial Intelligence. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Goldfarb, Avi, and Daniel Trefler. 2018. Artificial Intelligence and Trade. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Goldfarb, Avi, and Catherine Tucker. 2012. Privacy and Innovation. In *Innovation Policy and the Economy*. Volume 12. Eds. Josh Lerner and Scott Stern. NBER, University of Chicago Press, 65-89.

Goldfarb, Avi, and Catherine Tucker. 2017. Inequality, Privacy and Digital Market Design. Chapter in *Fair by Design*. Eds. Scott Kominers and Alex Teytelboym, Oxford University Press.

Goldfarb, Avi, and Catherine Tucker. Forthcoming. Digital Economics. *Journal of Economic Literature*.

Green, J. and Scotchmer, S. (1995) "On the division of profit in sequential innovation," *The RAND Journal of Economics* 26: 20-33.

Goldin, Claudia, and Lawrence F. Katz. 2010. *The Race Between Education and Technology*. Harvard University Press, Cambridge MA.

Hoffman, Mitchell; Lisa B. Kahn; and Danielle Li. 2018. Discretion in Hiring. *Quarterly Journal of Economics*. 133(2), 765-800.

Igami, Mitsuru. 2018. Artificial Intelligence as Structural Estimation: Economic Interpretations of Deep Blue, Bonanza, and AlphaGo. Working Paper, Yale University.

Imbens, Guido, and Donald Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press, Cambridge UK.

Kahneman, Daniel. 2018. Discussion of Camerer. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Lambrecht, Anja, and Catherine Tucker (2017). Algorithmic discrimination? Apparent algorithmic bias in the serving of stem ads. Mimeo, MIT.

Markov, John. 2015. *Machines of Loving Grace*. HarperCollins Publishers, New York.

Miller, A. and C. Tucker. 2018. Historic patterns of racial oppression and algorithms. Mimeo, MIT.

Mokyr, Joel. 2017. The Past and Future of Innovation: Some Lessons from Economic History. Working Paper, Northwestern University.

Moser, Petra. 2005. How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World Fairs. *American Economic Review* 95(4), 1214-1236.

Piketty, Thomas. 2013. *Capital in the twenty-first century*. Harvard University Press, Cambridge MA.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55(5): 999-1033.

Scott Morton, Fiona and Carl Shapiro. 2016. "Patent Assertions: Are We Any Closer to Aligning Reward to Contribution." In *Innovation Policy and the Economy* Volume 16. Eds. Josh Lerner and Scott Stern. P. 89-134. University of Chicago Press, Chicago.

Solow, Robert. 1987. We'd Better Watch Out. *New York Times Book Review*, July 12: 36.

Stevenson, Betsey. 2018. AI, Income, Employment, and Meaning. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Taddy, Matt. 2018. The Technological Elements of Artificial Intelligence. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Trajtenberg, Manuel. 2018. AI as the next GPT: A Political Economy Perspective. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Varian, Hal. 2018. Artificial Intelligence, Economics, and Industrial Organization. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Vincent, James. 2017. "Putin Says the Nation That Leads in AI 'Will Be the Ruler of the World,'" *The Verge*, September 4, 2017, <https://www.theverge.com/2017/9/4/16251226/russia-ai-putin-rule-the-world>.

Williams, Heidi. 2016. "Intellectual Property Rights and Innovation: Evidence from Health Care Markets." In *Innovation Policy and the Economy* Volume 16. Eds. Josh Lerner and Scott Stern. P. 53-88. University of Chicago Press, Chicago.