

The Impact of Digitalization on the Economy: A Review Article on the NBER Volume *Economics of Artificial Intelligence: An Agenda*

Eric Santor¹
Bank of Canada

ABSTRACT

Digitalization is affecting every aspect of our economy and our society. A set of new technologies are behind this latest surge - robotic process automation, artificial intelligence (AI) and machine learning (ML), big data, cloud computing, the internet of things and blockchain. This volume, *The Economics of Artificial Intelligence*, focuses on the impact, real and prospective, of machine learning (ML), on the economy. The authors tackle a wide range of topics, including how it is impacting innovation, the consequences for employment and economic growth, issues related to privacy, international trade and ultimately, how AI will affect the economics discipline itself. The contributors, overall, take a positive view of the impact of AI on economic outcomes. They also acknowledge, however, that policies related to redistribution, privacy and competition are needed to ensure that the benefits of digitalization are shared appropriately.

Digitalization is affecting every aspect of our economy and our society. A set of new technologies are behind this latest surge – robotic process automation, artificial intelligence (AI) and machine learning (ML), big data, cloud computing, the internet of things and blockchain. As these technologies became commercially viable, businesses are able to provide new products and

services, while reshaping how they operate. This volume, *The Economics of Artificial Intelligence: An Agenda*² published in 2019 by the University of Chicago as part of the National Bureau of Economic Research research volume series and edited by Ajay Agarwal, Joshua Gans and Avi Goldfarb, all at the University of Toronto, focuses on the impact, real and prospective, of arti-

¹ Eric Santor is Advisor to the Governor on Digitalization at the Bank of Canada. The views expressed here are those of the author and not those of the Governing Council. Email: esantor@bank-banque-canada.ca

² ISBN: 978-0-226-61333-8, USD\$130.00. The table of contents is available at <https://press.uchicago.edu/ucp/books/book/chicago/E/bo35780726.html>.

cial intelligence (AI), and specifically machine learning (ML), on the economy. The authors tackle a wide range of topics, including defining artificial intelligence and machine learning, how it is impacting innovation, and the consequences for employment and economic growth. The volume also explores important issues related to privacy, international trade and ultimately, how AI will affect the economics discipline itself.

The grand debate is whether AI will be a large net benefit for the economy, and ultimately, to society. Will ongoing digitalization lead us to a world more like Star Trek or Bladerunner? The contributors, overall, take a positive view of the impact of AI on economic outcomes. They also acknowledge that policies are needed to ensure that the benefits of digitalization are shared appropriately. It should be noted that some of the authors use AI and ML interchangeably.

The review unfolds as follows. The first section discusses whether AI/ML can be considered a General Purpose Technology and its importance for innovation. From there, the second section tackles the main issues related to the impact of ML on economic growth, employment and inequality. Policy issues are addressed in the third section, including privacy and competition. The fourth section looks at the implications for the discipline of economics. The review concludes with some suggestions for further investigation – specifically, competition and the rise of super-star firms.

Artificial Intelligence as a General Purpose Technology

The volume begins with an overarching summary of the existing debate on the impact of AI on economic growth by Erik Brynjolfsson, Daniel Rock and Chad Syverson. Their chapter, entitled “Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics”, explores why artificial intelligence is so important for economic growth, and the paradox that currently exists. The facts are well-known. Labour productivity growth slowed dramatically from 3 per cent annually in the 1950s and 1960s, to just over 1 per cent through the 1970s and 1980s. It picked up again with the advent of the internet and further gains in computing power from roughly 1995 to 2005. However, labour productivity has since slowed again to about 1 per cent. This slowdown has occurred even though a new range of technologies, and especially artificial intelligence, has burst onto the scene.

The authors highlight this simple paradox – if artificial intelligence is everywhere and a societal game-changer, why isn't it showing up in the productivity data? This play on Solow's famous quote on productivity is not unintentional. The authors then offer some possible explanations. First, while AI is very important for some sectors in the economy, it just might not deliver the expected productivity gains elsewhere. That said, the authors remain optimistic that AI will prove out as an economy-wide game-changer.

Second, the issue of mismeasurement is raised. It is widely recognized that official statistics are not able to keep up

with advances of technology and so actual productivity growth may be much higher than measured productivity. Nevertheless, the authors point to a series of studies to show that measurement bias cannot explain the productivity slowdown. In this regard, the volume in my view pays inadequate attention to the measurement issue. There are three shortcomings that stand out. First, there is evidence that while not initially measured, the productivity gains of the 1990s were only truly appreciated after historical revisions (Poloz, 2019). This may also be happening in the current context. Second, more attention is needed on highlighting exactly where the shortfalls in measurement methodologies or techniques are likely to be occurring. For example, official statistics are slow to include investments in data, data bases and data science, and in quality adjusting many of the digital services we now consume voraciously. Third, the authors note that the market power of large firms may be the culprit (as monopoly reduces the quantities supplied), but again remain unconvinced.

Instead, the most likely explanation is that, as a “General Purpose Technology” (GPT), it simply takes time for a technology like AI to diffuse through the economy. At the firm level, transformational GPTs require not only introducing the new tech, but also a completely new way of running a business. Consequently, it takes time, especially for incumbents, to adopt to the new ways of working. In the meantime, measured productivity will be low as investments are made and returns are initially slow to arrive. They point to the experience of previous GPTs to reinforce this point. Importantly, they make the simple

empirical point that current and past productivity trends do not predict future productivity growth.

The section then dives more deeply into what exactly is Artificial Intelligence and why it matters for innovation. Matt Taddy’s “The Technological Elements of Artificial Intelligence” provides a clear explanation of what is meant by AI. AI is not what is typically imagined in movies and popular conscience - machines thinking and acting like humans. AI is much more limited in the context of this volume to machine learning. Taddy usefully divides ML into General Purpose Machine Learning, Deep Learning and Reinforcement Learning. In each, a technical explanation is provided, but for the average reader the technical details are at times inaccessible. Greater use of concrete examples of how each type of technology is being used commercially would have been helpful.

The focus then turns to a simple question – if AI can lead to making better predictions at a lower cost, what role is there for humans in the process of decision-making? Ajay Agarwal, Joshua Gans and Avi Goldfarb’s “Prediction, Judgement and Complexity: A Theory of Decision-Making and Artificial Intelligence” sets out a model to show that, in fact, humans may be even more important in an AI-rich world. The model is simple: if the cost of prediction falls, then the value of its complements must rise. They argue that judgement is a complement and so demand for it will rise commensurately. How we can ensure that we bring the correct amount of judgement to an issue remains a topic for future research.

The rest of the section explores the im-

part of AI on innovation. Iain Cockburn, Rebecca Henderson and Scott Stern, in their chapter “The Impact of Artificial Intelligence on Innovation: an Exploratory Analysis”, show that AI, and especially machine learning, is likely to have a wide-ranging impact on how innovation is done. This is because it has the potential to be an “invention of a method of a method of invention” (Griliches, 1957). Specifically, AI-based learning will lower the cost of innovation by “automating discovery” and expanding the set of research and innovation opportunities that can be considered. They go on to document the incidence of AI in research output across disciplines and countries. Important lessons are also drawn - to make full use of ML, access to data is key, as well as appropriate competition and intellectual property policies. How AI can improve innovation is taken up by Ajay Agarwal, John McHale and Alexander Oettl in the follow-on chapter, “Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth”. In their view, AI can help to overcome the “burden of knowledge.” AI augments the ability of the researcher to combine new and existing knowledge in ways that were simply not possible before.

The last chapter of the section, “Artificial Intelligence as the Next GPT: A Political-Economy Perspective” offers thoughts on the political economy of technological disruption from AI. Manuel Trajtenberg focuses on the notion that there will be winners and losers of the coming transformation, and that governments will need to take appropriate action to support those workers who are displaced. Specifically, more focus will need to be paid to

soft skills, above and beyond the standard need for STEM skills, if the employment prospects of the displaced are to be enhanced. While providing some useful analysis, it is not clear how this chapter fits within the theme of the first section, which is AI as a GPT, and would more naturally be part of the next section.

Growth, Jobs and Inequality

When most people talk about the coming impact of AI on the economy, most of the attention is naturally focused on jobs. That is, will robots render humans superfluous in the production process, leading to wide-spread job losses and the societal ills that accompany such disruption? At the very least, ongoing automation of routine manual tasks, and the more recent automation of routine cognitive tasks, will exacerbate existing trends of stagnant real wages for many workers and the increase in inequality. The volume presents a wide range of views on various aspects of the impact of technological progress on employment, inequality and society.

Who is right? Are we just a few more advances in computing power away from a world of abundance and leisure, or is a dystopia of concentrated wealth amid widespread impoverishment? As a group, economists tend to have a positive view of technological progress, as highlighted by Betsey Stevenson’s chapter “Artificial Intelligence, Income, Employment and Meaning”. Over the long run, technological progress has led to higher incomes and fewer hours worked (allowing more time for leisure). But most discussions tend to focus on the short-run disruption – those neg-

actively affected by such technological advance.

The impact on employment is taken up more explicitly in the chapter “Artificial Intelligence, Automation, and Work” by Daron Acemoglu and Pascual Restrepo. They set out a model in which there are displacement effects and countervailing forces. Displacement effects are straightforward – technological advances lead to a reduction in the number of workers needed to do a task. The countervailing effects lean against the displacement effect. There are three main effects to consider: i) a productivity effect which increases the demand for non-automated tasks; ii) a capital accumulation effect which increases the demand for labour; and iii) a creation of new tasks effect – technological progress allows the production of new goods and services that were hitherto not possible nor even imagined. They usefully point out that the history of automation suggests that the third or “reinstatement effect” is quite powerful.

The relatively benign view of technological progress is echoed in the chapter by Phillipe Aghion, Benjamin Jones and Charles Jones “Artificial Intelligence and Economic Growth”. Setting out a simple growth model that incorporates Baumol’s cost effect and AI in the production function, they show that a balanced growth path can be achieved. That is, even though automation means faster productivity growth in the automated sector, this relative performance to the non-automated sector means that the automated sector could end up being a smaller share of the economy (due to more rapidly falling prices). In fact, long-run growth is determined by the sectors that are essential but

hard to improve with technology.

James Bessen takes a complementary approach in “Artificial Intelligence and Jobs: the Role of Demand”. Too often, pessimists focus on the income effects of technological progress and ignore the demand effects. Turning to historical evidence, he shows that the benefits from declines in prices due to higher productivity can far outweigh the negative effect on income. Using examples from the cotton, steel and automobile sectors, he shows that the decline in prices led to enormous increases in demand for those goods. AI-induced technological progress could be no different, and so we should not discount the ability of this sector to also create large demand effects.

Austan Goolsbee also takes a positive view of the impact of technological change on the labour market in “Public Policy in an AI Economy”. One of the key concerns of the pessimists is that the speed of adoption of AI may simply be faster than previous periods of technological advance, and as such, the displacement effects will be hard to offset. Using data from the Jobs Opening and Labor Turnover Survey (JOLTS), Goolsbee shows that one must put such potential job losses in context. If all 3.5 million jobs driving trucks, buses and cars were automated over the next 15 years, that would entail 19 000 job losses per month. This compares to average monthly job creation in the United States in 2017 of 5.3 million jobs (and 5.1 million separations). It seems the impact, from a macroeconomy point of view, will be less dramatic than many conceive.

Jason Furman, in his chapter “Should We Be Reassured If Automation in the Future Looks Like Automation in the Past?”,

also tackles the issue of labour market disruption, but with a focus on income inequality. On this issue, there is more room to be concerned. Both Goolsbee and Furman note that one solution to mitigate the impact on workers is a Universal Basic Income. While attractive on first glance, both raise concerns that will likely impede any immediate adoption. That said, the ongoing COVID-19 pandemic moves this issue to the fore as to how best to support incomes in a digital, post-pandemic world.

Jeffery Sachs takes a less optimistic view in his chapter “R&D, Structural Transformation, and the Distribution of Income”. He notes that the decline in the labour share of value added in the goods sector, from more than 60 per cent to just below 50 per cent, is evidence of how automation can affect the distribution of income. As AI-based technologies are adopted, the same phenomenon may appear for many medium-skilled workers in other sectors of the economy, with the gains accruing to the owners of capital. As such, one could expect income inequality to worsen. While Sachs provides interesting qualitative results, the quantitative impact remains to be done.

The theme of inequality is picked up more substantively in Anton Korinek and Joseph Stiglitz’s chapter “Artificial Intelligence and Its Implications for Income Distribution and Unemployment”. They posit that while there are concerns about labour substitution, the biggest issue will be one of income distribution. Technological change leads to increased inequality through two channels. The first channel is that technological change creates rents for the innovator at the expense of consumers. The sec-

ond channel is that AI may reduce wages for the types of workers who are being displaced to the benefit of entrepreneurs. To rectify these outcomes, they propose several policies – namely taxes on the rents and robust anti-trust policies. At the same time, workers that are negatively affected by automation can be compensated, either in the form of wage subsidies or earned income tax credits. Higher wages for public sector workers would also help lift wages for low-income earners and public investment could help to ensure higher aggregate demand.

Building on the overall pessimistic view of the impact of AI, the authors also consider the case of human enhancement and independent artificial intelligence development (that is, machines that make even better machines). Not surprisingly, both outcomes lead to even greater inequality, with the latter case even ending in some proportion of the population being unable to feed itself. While this is an interesting thought-experiment, its place in a practical volume such as this one is somewhat misplaced.

The section is rounded out by Tyler Cowen’s chapter “Neglected Open Questions in the Economics of Artificial Intelligence”. There are three areas of focus. First, he asks how consumer surplus will be divided. If the costs of the AI-enabled outputs are declining (as is typically the case for software-intensive goods), then it bodes well in terms of consumer surplus and displacement effects. This is because the benefits can be widely spread while the impact of the workers displaced in the production of the good will be limited. A second neglected issue are the international ef-

fects. These include how wage differentials will evolve in a world of even more pervasive AI-enabled outsourcing and whether it will be possible for developing countries to industrialize in the same manner as some emerging market economies have done in the past. Lastly, Cowen asks how political economy considerations will assist (or not) the process by which income can be redistributed effectively, e.g. to caring for elderly persons. Needless to say, much research remains to be done to address these and many other questions.

Machine Learning and Regulation

This section moves the discussion to how AI might affect firms' organizational structure, privacy issues, and commensurately, the implications for regulation of firms.

Hal Varian's chapter "Artificial Intelligence, Economics and Industrial Organization" provides a wide-ranging overview of how the adoption of AI will affect businesses' organizational structure. The chapter begins with a summary of what machine learning can do but adds little new given the previous chapters. The more interesting contribution comes in the form of how firms themselves may change. For instance, it is not immediately obvious if ML will increase or decrease the minimum efficient scale for a firm. Likewise, price discrimination can be more easily implemented with the use of ML and big data, but to what extent firms will aim for perfect price discrimination is unknown. Varian also raises some key concerns, such as collusion, security and privacy, which are tackled more comprehensively in the rest of the section.

While much of the discussion on the impact of AI on the economy is focused on economic issues, it is privacy that often receives the most attention by the wider public. Firms' use of ML to make better predictions and affect consumer behaviour requires large and detailed data on consumers. How these data are managed and used is often "front-page" news. Two chapters make very useful contributions on this front. The chapter "Privacy, Algorithms, and Artificial Intelligence" by Catherine Tucker outlines three key areas of focus: data persistence, data repurposing and data spillovers. Data persistence is a key attribute of the current environment. Data created through transactions, unlike in previous times, are very likely to persist. This creates a time-consistency problem. Data that a person may have felt comfortable giving when they were younger may not necessarily be that which their older self may find acceptable. While there is evidence that some of the biggest players in the ML space do not need long data histories to do their work, some forms of data may be very persistent. For example, genetic data can be very long lived (a lifetime or even beyond) and so how it is used and controlled is of considerable import.

The second issue is data repurposing. Data collected for one purpose may end up being used in a completely different context, often without the knowledge of the consumer. This lack of knowledge of use can be damaging to consumers. A third issue is that data spillovers are a natural outcome of ML algorithms themselves. The most concerning of these spillovers is the occurrence of algorithmic bias or discrimination. Recent high-profile cases have

shown that ML algorithms can introduce or reinforce existing biases if left unattended. For example, Amazon stopped an ML algorithm used for recruitment as it was seen to be biased against women. This reinforces the need for firms to use their data in an ethical manner.

The chapter “Artificial Intelligence and Consumer Privacy” by Ginger Zhe Jin highlights some of the new risks associated with the implementation of ML and big data in the consumer context. At the heart of these risks is that the firm may promise to respect data privacy *ex ante*, but renege on the promise *ex post*. Moreover, it is very hard for the consumer to even know if the promise of data privacy is being respected. High-profile breach cases, such as Target and Equifax, highlight the risks to both consumers and firms if they fail to keep their data secure. Firms may lose the trust of their customers, and customers may become the targets of misuse, ranging from fraud, identity theft, and even blackmail or extortion. Even if data are kept secure, Jin notes that data may not be restricted to their intended use, potentially harming the customer.

To reduce these risks, Jin points to two sources of mitigation: consumer attitudes and supply-side actions. Consumer attitudes are not a panacea. While many consumers profess concerns about privacy, in practice they are less strict in how they control their data. While this seems a paradox, it could be the case that consumers are willing to take the risk of giving away their data given the benefits they might receive from firms knowing them well. The more likely solution lies in supply-side actions. Firms actions have been mixed.

While many firms fear the market and reputational consequences of data breaches and invest in data security accordingly, the average firm suffers little material consequence when a breach occurs. Consequently, it will be the responsibility of the government to appropriately regulate and enforce data privacy and appropriate use. On this front, however, there is much to do. In the United States, for example, there is a patchwork of regulation and federal law in which firms and consumers operate. Shoring up this system can be done in many ways – safety regulation, consumer notice and disclosure and self-regulation – but each will need to be mindful of not thwarting the innovation that ultimately benefits consumers.

The question of who should be punished if AI goes astray, and the implications for innovation, is tackled in the contribution by Alberto Galasso and Hong Luo, “Punishing Robots: Issues to the Economics of Tort Liability and Innovation in Artificial Intelligence”. They present a simple model of the impact of liability risk on the incentives to innovate. The bottom-line is that liability risk does not necessarily chill innovation – rather, there is a complex, and often ambiguous relationship that depends on the expected payoffs of innovation relative to its liability risk profile. They then go on to enumerate some of the key liability issues as they relate to AI technology. Assigning liability to AI that goes awry is not so straightforward, and crafting a suitable legal framework that does not inhibit innovation remains an open question.

The last chapter in the section addresses the international dimensions of AI, an oft neglected subject. Avi Goldfarb and Daniel

Trefler, in “Artificial Intelligence and International Trade” focus on some elements that characterize AI – the benefits that come from economies of scope and scale, and knowledge externalities. As such, trade policies need to take account of these features, being mindful that some policies that are aimed at consumers (such as privacy policies, data localization) can end up as disguised trade polices.

Machine Learning and Economics

The last section of the book focuses on the implications of ML on the discipline of economics. The first contribution by Susan Athey, “The Impact of Machine Learning on Economics” highlights the many ways in which machine learning can enhance econometric analysis. By leveraging big data and its inherently iterative approach, ML has the potential to improve the forecasting performance of key economic variables, and in doing so, allowing for better policy action. It will also likely allow for the investigation of economic questions that were not feasible in the past. In this context, some more concrete examples would be helpful along with the link to what it might entail for policy. Athey also correctly notes the current limitations of ML with respect to causal inference. While much progress is being made in improving the “explainability” of ML algorithms, there is still much work to be done before policymakers will be at ease with ML-based results.

The impact of ML on economics will also depend on the availability of data. In “Artificial Intelligence, Labor, Productivity, and the Need for Firm-Level Data”,

Manav Raj and Robert Seamans highlight the need for more and better data at the firm level if we are to better understand how AI is going to affect the economy. While there is some progress on this front, large data gaps remain. They offer some solutions, including a census of robots and surveys at the firm level. Unfortunately, for researchers, much firm-level data are proprietary, and will likely remain inaccessible in the near-term. In this sense, deeper collaboration between researchers and the private sector could be pursued.

The last two chapters look at the implications of AI/ML for market design and behavioural economics. In both cases, the insight is simple. ML provides the opportunity to better understand the economic phenomena at hand to inform decision-making. In the case of market design, Paul Milgrom and Steven Tadelis in “How Artificial Intelligence and Machine Learning Can Impact Market Design” showcase how ML was used to consider the implications of re-allocating frequency bandwidth in the US market. Simply, without ML, it would have been almost impossible to consider the set of options at hand in a systematic fashion. Colin Camerer, in “Artificial Intelligence and Behavioural Economics”, shows how ML could be leveraged to better model behavioural choices. This is because ML allows the researcher to consider a wider range of factors that could affect decision making than more standard models.

Conclusion

The volume provides a comprehensive view of the economics of artificial intelligence. The volume is for the most part accessible, as most of the contributions can be understood with a basic level of economics. However, from time to time, some of the chapters present models that would be beyond the comprehension of many intended readers.

A wide range of topics are covered in the book, with an appropriate focus on why AI could lead to (or not) a pickup in productivity growth and the expected impacts on the labour market. In this respect, the volume successfully presents both sides of the debate in a constructive manner. However, it might have been useful to have a) some discussion of what happens if the anticipated productivity gains do not occur and b) a summary chapter on the issue of whether AI boosts trend productivity

growth in order to form an overall view on the subject.

There are two issues that only received cursory attention. The issue of the impact of AI on competition could be better explored. The rise of superstar firms, driven by AI, has raised many anti-trust issues and increases in industry concentration is seen as one of the main contributors to ongoing inequality. Relatedly, some discussion of who owns and controls data would have been insightful. That said, for an informed reader, this volume provides a great primer on the economics of artificial intelligence.

References

- Griliches, Z. (1957) "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica*, Vol. 25, No. 4, pp. 501-22.
- Poloz, Stephen S. (2019) "Technological Progress and Monetary Policy: Managing the Fourth Industrial Revolution," Bank of Canada Staff Discussion Paper No. 2019-11.